1. Ice Manual . . . 15
1.1 Ice Overview . . . 17
1.1.1 Ice Architecture . . . 18
1.1.1.1 Terminology . . . 19
1.1.1.2 Slice (Specification Language for Ice) . . . 26
1.1.1.3 Overview of the Language Mappings . . . 27
1.1.1.4 Client and Server Structure . . . 28
1.1.1.5 Overview of the Ice Protocol . . . 30
1.1.2 Ice Services Overview . . . 31
1.1.3 Architectural Benefits of Ice . . . 33
1.2 Hello World Application . . . 34
1.2.1 Writing a Slice Definition . . . 35
1.2.2 Writing an Ice Application with C++ . . . 36
1.2.3 Writing an Ice Application with C-Sharp . . . 43
1.2.4 Writing an Ice Application with Java . . . 50
1.2.5 Writing an Ice Application with JavaScript . . . 57
1.2.6 Writing an Ice Application with Objective-C . . . 60
1.2.7 Writing an Ice Application with PHP . . . 67
1.2.8 Writing an Ice Application with Python . . . 70
1.2.9 Writing an Ice Application with Ruby . . . 75
1.2.10 Writing an Ice Application with Visual Basic . . . 78
1.3 The Slice Language . . . 85
1.3.1 Slice Compilation . . . 86
1.3.2 Slice Source Files . . . 89
1.3.3 Lexical Rules . . . 92
1.3.4 Modules . . . 94
1.3.5 Basic Types . . . 96
1.3.6 User-Defined Types . . . 98
1.3.6.1 Enumerations . . . 99
1.3.6.2 Structures . . . 101
1.3.6.3 Sequences . . . 103
1.3.6.4 Dictionaries . . . 105
1.3.7 Constants and Literals . . . 107
1.3.8 Interfaces, Operations, and Exceptions . . . 111
1.3.8.1 Operations . . . 112
1.3.8.2 User Exceptions . . . 117
1.3.8.3 Run-Time Exceptions . . . 122
1.3.8.4 Proxies for Ice Objects . . . 126
1.3.8.5 Interface Inheritance . . . 128
1.3.9 Classes . . . 136
1.3.9.1 Simple Classes . . . 137
1.3.9.2 Class Inheritance . . . 139
1.3.9.3 Class Inheritance Semantics . . . 141
1.3.9.4 Classes as Unions . . . 143
1.3.9.5 Self-Referential Classes . . . 144
1.3.9.6 Classes Versus Structures . . . 147
1.3.9.7 Classes with Operations . . . 148
1.3.9.8 Architectural Implications of Classes . . . 149
1.3.9.9 Classes Implementing Interfaces . . . 151
1.3.9.10 Class Inheritance Limitations . . . 154
1.3.9.11 Pass-by-Value Versus Pass-by-Reference . . . 155
1.3.9.12 Passing Interfaces by Value . . . 157
1.3.9.13 Classes with Compact Type IDs . . . 158
1.3.10 Forward Declarations . . . 159
1.3.11 Optional Data Members . . . 160
1.3.12 Type IDs . . . 163
1.3.13 Operations on Object . . . 164
1.3.14 Local Types . . . 166
1.3.15 Names and Scoping . . . 167
1.3.16 Metadata . . . 177
1.3.17 Serializable Objects . . . 178
1.3.18 Deprecating Slice Definitions . . . 180
1.3.19 Using the Slice Compilers . . . 181
1.3.20 Slice Checksums . . . 183
1.3.21 Generating Slice Documentation . . . 184
1.3.22 Slice Keywords . . . 190
1.3.23 Slice Metadata Directives . . . 191
1.3.24 Slice for a Simple File System . . . 198
1.4 Language Mappings . . . 201
1.4.1 C++ Mapping . . . 202
1.4.1.1 Client-Side Slice-to-C++ Mapping . . . 203
1.4.1.1.1 C++ Mapping for Identifiers . . . 204

Copyright 2016, ZeroC, Inc.

1.4.1.1.2 C++ Mapping for Modules . . . 205
1.4.1.1.3 C++ Mapping for Built-In Types . . . 206
1.4.1.1.4 C++ Mapping for Enumerations . . . 209
1.4.1.1.5 C++ Mapping for Structures . . . 210
1.4.1.1.6 C++ Mapping for Sequences . . . 214
1.4.1.1.7 C++ Mapping for Dictionaries . . . 218
1.4.1.1.8 C++ Mapping for Constants . . . 220
1.4.1.1.9 C++ Mapping for Exceptions . . . 222
1.4.1.1.10 C++ Mapping for Interfaces . . . 227
1.4.1.1.11 C++ Mapping for Operations . . . 236
1.4.1.1.12 C++ Mapping for Optional Values . . . 246
1.4.1.1.13 C++ Mapping for Classes . . . 248
1.4.1.1.14 Smart Pointers for Classes . . . 257
1.4.1.1.15 Asynchronous Method Invocation (AMI) in C++ . . . 271
1.4.1.1.16 Using Slice Checksums in C++ . . . 286
1.4.1.1.17 Example of a File System Client in C++ . . . 287
1.4.1.2 Server-Side Slice-to-C++ Mapping . . . 291
1.4.1.2.1 The Server-Side main Function in C++ . . . 292
1.4.1.2.2 Server-Side C++ Mapping for Interfaces . . . 304
1.4.1.2.3 Parameter Passing in C++ . . . 307
1.4.1.2.4 Raising Exceptions in C++ . . . 309
1.4.1.2.5 Object Incarnation in C++ . . . 310
1.4.1.2.6 Asynchronous Method Dispatch (AMD) in C++ . . . 315
1.4.1.2.7 Example of a File System Server in C++ . . . 320
1.4.1.3 Customizing the C++ Mapping . . . 336
1.4.1.3.1 The C++ Stream Helpers . . . 337
1.4.1.3.2 The cpp:type and cpp:view-type Metadata Directives . . . 345
1.4.1.4 Version Information in C++ . . . 358
1.4.1.5 slice2cpp Command-Line Options . . . 359
1.4.1.6 C++ Strings and Character Encoding . . . 365
1.4.1.6.1 Installing String Converters . . . 366
1.4.1.6.2 UTF-8 Conversion . . . 368
1.4.1.6.3 String Parameters in Local Calls . . . 369
1.4.1.6.4 Built-in String Converters . . . 370
1.4.1.6.5 String Conversion Convenience Functions . . . 371
1.4.1.6.6 The iconv String Converter . . . 373
1.4.1.6.7 The Ice String Converter Plug-in . . . 374
1.4.1.6.8 Custom String Converter Plug-ins . . . 375
1.4.1.7 The C++ Utility Library . . . 376
1.4.1.7.1 Threads and Concurrency with C++ . . . 377
1.4.1.7.2 The C++ AbstractMutex Class . . . 409
1.4.1.7.3 The C++ Cache Template . . . 412
1.4.1.7.4 The C++ Exception Class . . . 414
1.4.1.7.5 The C++ generateUUID Function . . . 415
1.4.1.7.6 The C++ Handle Template . . . 416
1.4.1.7.7 The C++ Handle Template Adaptors . . . 420
1.4.1.7.8 The C++ ScopedArray Template . . . 426
1.4.1.7.9 The C++ Shared and SimpleShared Classes . . . 428
1.4.1.7.10 The C++ Time Class . . . 429
1.4.1.7.11 The C++ Timer and TimerTask Classes . . . 433
1.4.2 C-Sharp Mapping . . . 435
1.4.2.1 Client-Side Slice-to-C-Sharp Mapping . . . 436
1.4.2.1.1 C-Sharp Mapping for Identifiers . . . 437
1.4.2.1.2 C-Sharp Mapping for Modules . . . 438
1.4.2.1.3 C-Sharp Mapping for Built-In Types . . . 439
1.4.2.1.4 C-Sharp Mapping for Enumerations . . . 440
1.4.2.1.5 C-Sharp Mapping for Structures . . . 441
1.4.2.1.6 C-Sharp Mapping for Sequences . . . 447
1.4.2.1.7 C-Sharp Mapping for Dictionaries . . . 451
1.4.2.1.8 C-Sharp Collection Comparison . . . 453
1.4.2.1.9 C-Sharp Mapping for Constants . . . 454
1.4.2.1.10 C-Sharp Mapping for Exceptions . . . 456
1.4.2.1.11 C-Sharp Mapping for Interfaces . . . 460
1.4.2.1.12 C-Sharp Mapping for Operations . . . 468
1.4.2.1.13 C-Sharp Mapping for Classes . . . 477
1.4.2.1.14 C-Sharp Mapping for Optional Values . . . 488
1.4.2.1.15 Serializable Objects in C-Sharp . . . 490
1.4.2.1.16 C-Sharp Specific Metadata Directives . . . 492
1.4.2.1.17 Asynchronous Method Invocation (AMI) in C-Sharp . . . 493
1.4.2.1.18 slice2cs Command-Line Options . . . 507
1.4.2.1.19 Using Slice Checksums in C-Sharp . . . 508
1.4.2.1.20 Example of a File System Client in C-Sharp . . . 509

Copyright 2016, ZeroC, Inc.

1.4.2.2 Server-Side Slice-to-C-Sharp Mapping . . . 513
1.4.2.2.1 The Server-Side main Method in C-Sharp . . . 514
1.4.2.2.2 Server-Side C-Sharp Mapping for Interfaces . . . 521
1.4.2.2.3 Parameter Passing in C-Sharp . . . 525
1.4.2.2.4 Raising Exceptions in C-Sharp . . . 527
1.4.2.2.5 Tie Classes in C-Sharp . . . 528
1.4.2.2.6 Object Incarnation in C-Sharp . . . 532
1.4.2.2.7 Asynchronous Method Dispatch (AMD) in C-Sharp . . . 536
1.4.2.2.8 Example of a File System Server in C-Sharp . . . 541
1.4.2.3 The .NET Utility Library . . . 551
1.4.3 Java Mapping . . . 553
1.4.3.1 Client-Side Slice-to-Java Mapping . . . 554
1.4.3.1.1 Java Mapping for Identifiers . . . 555
1.4.3.1.2 Java Mapping for Modules . . . 556
1.4.3.1.3 Java Mapping for Built-In Types . . . 557
1.4.3.1.4 Java Mapping for Enumerations . . . 558
1.4.3.1.5 Java Mapping for Structures . . . 560
1.4.3.1.6 Java Mapping for Sequences . . . 563
1.4.3.1.7 Java Mapping for Dictionaries . . . 564
1.4.3.1.8 Java Mapping for Constants . . . 565
1.4.3.1.9 Java Mapping for Exceptions . . . 567
1.4.3.1.10 Java Mapping for Interfaces . . . 571
1.4.3.1.11 Java Mapping for Operations . . . 579
1.4.3.1.12 Java Mapping for Classes . . . 588
1.4.3.1.13 Java Mapping for Optional Data Members . . . 596
1.4.3.1.14 Serializable Objects in Java . . . 599
1.4.3.1.15 Customizing the Java Mapping . . . 601
1.4.3.1.16 Asynchronous Method Invocation (AMI) in Java . . . 613
1.4.3.1.17 Using the Slice Compiler for Java . . . 626
1.4.3.1.18 slice2java Command-Line Options . . . 627
1.4.3.1.19 Using Slice Checksums in Java . . . 628
1.4.3.1.20 Example of a File System Client in Java . . . 629
1.4.3.2 Server-Side Slice-to-Java Mapping . . . 634
1.4.3.2.1 The Server-Side main Method in Java . . . 635
1.4.3.2.2 Server-Side Java Mapping for Interfaces . . . 641
1.4.3.2.3 Parameter Passing in Java . . . 645
1.4.3.2.4 Raising Exceptions in Java . . . 648
1.4.3.2.5 Tie Classes in Java . . . 650
1.4.3.2.6 Object Incarnation in Java . . . 654
1.4.3.2.7 Asynchronous Method Dispatch (AMD) in Java . . . 658
1.4.3.2.8 Example of a File System Server in Java . . . 663
1.4.3.3 The Java Utility Library . . . 673
1.4.3.4 Custom Class Loaders . . . 676
1.4.3.5 Java Interrupts . . . 677
1.4.4 JavaScript Mapping . . . 679
1.4.4.1 Asynchronous APIs in JavaScript . . . 681
1.4.4.2 Client-Side Slice-to-JavaScript Mapping . . . 682
1.4.4.2.1 JavaScript Mapping for Identifiers . . . 683
1.4.4.2.2 JavaScript Mapping for Modules . . . 684
1.4.4.2.3 JavaScript Mapping for Built-In Types . . . 686
1.4.4.2.4 JavaScript Mapping for Enumerations . . . 687
1.4.4.2.5 JavaScript Mapping for Structures . . . 689
1.4.4.2.6 JavaScript Mapping for Sequences . . . 691
1.4.4.2.7 JavaScript Mapping for Dictionaries . . . 692
1.4.4.2.8 JavaScript Mapping for Constants . . . 695
1.4.4.2.9 JavaScript Mapping for Exceptions . . . 697
1.4.4.2.10 JavaScript Mapping for Interfaces . . . 700
1.4.4.2.11 JavaScript Mapping for Operations . . . 707
1.4.4.2.12 JavaScript Mapping for Classes . . . 722
1.4.4.2.13 slice2js Command-Line Options . . . 729
1.4.4.3 Server-Side Slice-to-JavaScript Mapping . . . 730
1.4.4.3.1 Server-Side JavaScript Mapping for Interfaces . . . 731
1.4.4.3.2 Parameter Passing in JavaScript . . . 734
1.4.4.3.3 Raising Exceptions in JavaScript . . . 736
1.4.4.3.4 Object Incarnation in JavaScript . . . 737
1.4.4.3.5 Asynchronous Method Dispatch (AMD) in JavaScript . . . 741
1.4.5 Objective-C Mapping . . . 745
1.4.5.1 Client-Side Slice-to-Objective-C Mapping . . . 746
1.4.5.1.1 Objective-C Mapping for Modules . . . 747
1.4.5.1.2 Objective-C Mapping for Identifiers . . . 749
1.4.5.1.3 Objective-C Mapping for Built-In Types . . . 751
1.4.5.1.4 Objective-C Mapping for Enumerations . . . 752

Copyright 2016, ZeroC, Inc.

1.4.5.1.5 Objective-C Mapping for Structures . . . 753
1.4.5.1.6 Objective-C Mapping for Sequences . . . 757
1.4.5.1.7 Objective-C Mapping for Dictionaries . . . 761
1.4.5.1.8 Objective-C Mapping for Constants . . . 763
1.4.5.1.9 Objective-C Mapping for Exceptions . . . 765
1.4.5.1.10 Objective-C Mapping for Interfaces . . . 773
1.4.5.1.11 Objective-C Mapping for Operations . . . 778
1.4.5.1.12 Objective-C Mapping for Local Interfaces . . . 790
1.4.5.1.13 Objective-C Mapping for Classes . . . 791
1.4.5.1.14 Objective-C Mapping for Interfaces by Value . . . 801
1.4.5.1.15 Objective-C Mapping for Optional Data Members . . . 802
1.4.5.1.16 Asynchronous Method Invocation (AMI) in Objective-C . . . 803
1.4.5.1.17 slice2objc Command-Line Options . . . 812
1.4.5.1.18 Example of a File System Client in Objective-C . . . 813
1.4.5.2 Server-Side Slice-to-Objective-C Mapping . . . 818
1.4.5.2.1 The Server-Side main Function in Objective-C . . . 819
1.4.5.2.2 Server-Side Objective-C Mapping for Interfaces . . . 822
1.4.5.2.3 Parameter Passing in Objective-C . . . 826
1.4.5.2.4 Raising Exceptions in Objective-C . . . 828
1.4.5.2.5 Object Incarnation in Objective-C . . . 829
1.4.5.2.6 Example of a File System Server in Objective-C . . . 833
1.4.6 PHP Mapping . . . 848
1.4.6.1 Client-Side Slice-to-PHP Mapping . . . 849
1.4.6.1.1 PHP Mapping for Identifiers . . . 850
1.4.6.1.2 PHP Mapping for Modules . . . 851
1.4.6.1.3 PHP Mapping for Built-In Types . . . 852
1.4.6.1.4 PHP Mapping for Enumerations . . . 853
1.4.6.1.5 PHP Mapping for Structures . . . 855
1.4.6.1.6 PHP Mapping for Sequences . . . 857
1.4.6.1.7 PHP Mapping for Dictionaries . . . 858
1.4.6.1.8 PHP Mapping for Constants . . . 859
1.4.6.1.9 PHP Mapping for Exceptions . . . 861
1.4.6.1.10 PHP Mapping for Interfaces . . . 865
1.4.6.1.11 PHP Mapping for Operations . . . 872
1.4.6.1.12 PHP Mapping for Classes . . . 879
1.4.6.1.13 slice2php Command-Line Options . . . 887
1.4.6.1.14 Application Notes for PHP . . . 888
1.4.6.1.15 Using Slice Checksums in PHP . . . 895
1.4.6.1.16 Example of a File System Client in PHP . . . 896
1.4.7 Python Mapping . . . 900
1.4.7.1 Client-Side Slice-to-Python Mapping . . . 901
1.4.7.1.1 Python Mapping for Identifiers . . . 902
1.4.7.1.2 Python Mapping for Modules . . . 903
1.4.7.1.3 Python Mapping for Built-In Types . . . 904
1.4.7.1.4 Python Mapping for Enumerations . . . 905
1.4.7.1.5 Python Mapping for Structures . . . 908
1.4.7.1.6 Python Mapping for Sequences . . . 910
1.4.7.1.7 Python Mapping for Dictionaries . . . 913
1.4.7.1.8 Python Mapping for Constants . . . 914
1.4.7.1.9 Python Mapping for Exceptions . . . 916
1.4.7.1.10 Python Mapping for Interfaces . . . 920
1.4.7.1.11 Python Mapping for Operations . . . 927
1.4.7.1.12 Python Mapping for Classes . . . 934
1.4.7.1.13 Asynchronous Method Invocation (AMI) in Python . . . 942
1.4.7.1.14 Code Generation in Python . . . 951
1.4.7.1.15 Using Slice Checksums in Python . . . 961
1.4.7.1.16 Example of a File System Client in Python . . . 962
1.4.7.2 Server-Side Slice-to-Python Mapping . . . 966
1.4.7.2.1 The Server-Side main Program in Python . . . 967
1.4.7.2.2 Server-Side Python Mapping for Interfaces . . . 972
1.4.7.2.3 Parameter Passing in Python . . . 975
1.4.7.2.4 Raising Exceptions in Python . . . 977
1.4.7.2.5 Object Incarnation in Python . . . 978
1.4.7.2.6 Asynchronous Method Dispatch (AMD) in Python . . . 982
1.4.7.2.7 Example of a File System Server in Python . . . 986
1.4.8 Ruby Mapping . . . 995
1.4.8.1 Client-Side Slice-to-Ruby Mapping . . . 996
1.4.8.1.1 Ruby Mapping for Identifiers . . . 997
1.4.8.1.2 Ruby Mapping for Modules . . . 998
1.4.8.1.3 Ruby Mapping for Built-In Types . . . 999
1.4.8.1.4 Ruby Mapping for Enumerations . . . 1000
1.4.8.1.5 Ruby Mapping for Structures . . . 1003

Copyright 2016, ZeroC, Inc.

1.4.8.1.6 Ruby Mapping for Sequences . . . 1005
1.4.8.1.7 Ruby Mapping for Dictionaries . . . 1007
1.4.8.1.8 Ruby Mapping for Constants . . . 1008
1.4.8.1.9 Ruby Mapping for Exceptions . . . 1010
1.4.8.1.10 Ruby Mapping for Interfaces . . . 1014
1.4.8.1.11 Ruby Mapping for Operations . . . 1021
1.4.8.1.12 Ruby Mapping for Classes . . . 1028
1.4.8.1.13 Code Generation in Ruby . . . 1037
1.4.8.1.14 The main Program in Ruby . . . 1043
1.4.8.1.15 Using Slice Checksums in Ruby . . . 1048
1.4.8.1.16 Example of a File System Client in Ruby . . . 1049
1.5 Properties and Configuration . . . 1054
1.5.1 Properties Overview . . . 1055
1.5.2 Configuration File Syntax . . . 1057
1.5.3 Setting Properties on the Command Line . . . 1060
1.5.4 Using Configuration Files . . . 1061
1.5.5 Alternate Property Stores . . . 1064
1.5.6 Command-Line Parsing and Initialization . . . 1065
1.5.7 The Properties Interface . . . 1068
1.5.8 Reading Properties . . . 1070
1.5.9 Setting Properties . . . 1072
1.5.10 Parsing Properties . . . 1077
1.6 Communicator and other Core Local Features . . . 1082
1.6.1 Communicators . . . 1083
1.6.2 Communicator Initialization . . . 1087
1.6.3 Object Identity . . . 1089
1.6.4 Plug-in Facility . . . 1093
1.6.4.1 Plug-in API . . . 1094
1.6.4.2 Plug-in Configuration . . . 1098
1.6.4.3 Advanced Plug-in Topics . . . 1099
1.7 Client-Side Features . . . 1101
1.7.1 Proxies . . . 1102
1.7.1.1 Obtaining Proxies . . . 1103
1.7.1.2 Proxy Methods . . . 1106
1.7.1.3 Proxy Endpoints . . . 1111
1.7.1.4 Filtering Proxy Endpoints . . . 1112
1.7.1.5 Proxy Defaults and Overrides . . . 1113
1.7.1.6 Proxy-Based Load Balancing . . . 1115
1.7.1.7 Proxy and Endpoint Syntax . . . 1117
1.7.2 Request Contexts . . . 1124
1.7.2.1 Explicit Request Contexts . . . 1125
1.7.2.2 Per-Proxy Request Contexts . . . 1127
1.7.2.3 Implicit Request Contexts . . . 1129
1.7.2.4 Design Considerations for Request Contexts . . . 1131
1.7.3 Invocation Timeouts . . . 1133
1.7.4 Automatic Retries . . . 1135
1.7.5 Oneway Invocations . . . 1139
1.7.6 Datagram Invocations . . . 1143
1.7.7 Batched Invocations . . . 1146
1.8 Server-Side Features . . . 1152
1.8.1 Object Adapters . . . 1153
1.8.1.1 The Active Servant Map . . . 1154
1.8.1.2 Creating an Object Adapter . . . 1156
1.8.1.3 Servant Activation and Deactivation . . . 1157
1.8.1.4 Object Adapter States . . . 1159
1.8.1.5 Object Adapter Endpoints . . . 1162
1.8.1.6 Creating Proxies with an Object Adapter . . . 1166
1.8.1.7 Using Multiple Object Adapters . . . 1168
1.8.2 The Current Object . . . 1169
1.8.3 Servant Locators . . . 1171
1.8.3.1 The ServantLocator Interface . . . 1172
1.8.3.2 Threading Guarantees for Servant Locators . . . 1174
1.8.3.3 Registering a Servant Locator . . . 1175
1.8.3.4 Servant Locator Example . . . 1177
1.8.3.5 Using Identity Categories with Servant Locators . . . 1181
1.8.3.6 Using Cookies with Servant Locators . . . 1184
1.8.4 Default Servants . . . 1185
1.8.5 Dispatch Interceptors . . . 1189
1.9 Client-Server Features . . . 1194
1.9.1 The Ice Threading Model . . . 1195
1.9.1.1 Thread Pools . . . 1196
1.9.1.2 Object Adapter Thread Pools . . . 1198

Copyright 2016, ZeroC, Inc.

1.9.1.3 Thread Pool Design Considerations . . . 1199
1.9.1.4 Concurrent Proxy Invocations . . . 1201
1.9.1.5 Nested Invocations . . . 1202
1.9.1.6 Thread Safety . . . 1204
1.9.1.7 Dispatching Invocations to User Threads . . . 1210
1.9.1.8 Blocking API Calls . . . 1225
1.9.2 Connection Management . . . 1226
1.9.2.1 Connection Establishment . . . 1227
1.9.2.2 Active Connection Management . . . 1231
1.9.2.3 Using Connections . . . 1235
1.9.2.4 Connection Closure . . . 1243
1.9.2.5 Bidirectional Connections . . . 1244
1.9.3 Connection Timeouts . . . 1248
1.9.4 Collocated Invocation and Dispatch . . . 1251
1.9.5 Locators . . . 1253
1.9.5.1 Locator Semantics for Clients . . . 1254
1.9.5.2 Locator Configuration for a Client . . . 1257
1.9.5.3 Locator Semantics for Servers . . . 1258
1.9.5.4 Locator Configuration for a Server . . . 1259
1.9.6 Slicing Values and Exceptions . . . 1261
1.9.7 Dynamic Ice . . . 1272
1.9.7.1 Streaming Interfaces . . . 1273
1.9.7.1.1 C++ Streaming Interfaces . . . 1274
1.9.7.1.2 Java Streaming Interfaces . . . 1292
1.9.7.1.3 C-Sharp Streaming Interfaces . . . 1310
1.9.7.2 Dynamic Invocation and Dispatch . . . 1329
1.9.7.2.1 Dynamic Invocation and Dispatch Overview . . . 1330
1.9.7.2.2 Dynamic Invocation and Dispatch in C++ . . . 1333
1.9.7.2.3 Dynamic Invocation and Dispatch in Java . . . 1338
1.9.7.2.4 Dynamic Invocation and Dispatch in C-Sharp . . . 1344
1.9.7.3 Asynchronous Dynamic Invocation and Dispatch . . . 1350
1.9.7.3.1 Asynchronous Dynamic Invocation and Dispatch in C++ . . . 1351
1.9.7.3.2 Asynchronous Dynamic Invocation and Dispatch in Java . . . 1357
1.9.7.3.3 Asynchronous Dynamic Invocation and Dispatch in C-Sharp . . . 1361
1.9.8 Facets . . . 1364
1.9.9 Versioning . . . 1371
1.9.9.1 Versioning through Incremental Updates . . . 1372
1.9.9.2 Versioning with Facets . . . 1375
1.10 Administration and Diagnostics . . . 1380
1.10.1 Administrative Facility . . . 1381
1.10.1.1 The admin Object . . . 1382
1.10.1.2 Creating the admin Object . . . 1383
1.10.1.3 Using the admin Object . . . 1384
1.10.1.4 The Process Facet . . . 1386
1.10.1.5 The Properties Facet . . . 1389
1.10.1.6 The Logger Facet . . . 1394
1.10.1.7 The Metrics Facet . . . 1398
1.10.1.8 Filtering Administrative Facets . . . 1403
1.10.1.9 Custom Administrative Facets . . . 1404
1.10.1.10 Security Considerations for Administrative Facets . . . 1405
1.10.2 Logger Facility . . . 1406
1.10.2.1 The Default Logger . . . 1407
1.10.2.2 Custom Loggers . . . 1408
1.10.2.3 Built-in Loggers . . . 1409
1.10.2.4 Logger Plug-ins . . . 1410
1.10.2.5 The Per-Process Logger . . . 1415
1.10.2.6 C++ Logger Utility Classes . . . 1416
1.10.3 Instrumentation Facility . . . 1418
1.11 IceBox . . . 1423
1.11.1 Developing IceBox Services . . . 1424
1.11.2 Configuring IceBox Services . . . 1430
1.11.3 Starting the IceBox Server . . . 1435
1.11.4 IceBox Administration . . . 1437
1.12 Ice Plugins . . . 1441
1.12.1 IceDiscovery . . . 1442
1.12.2 IceLocatorDiscovery . . . 1447
1.12.3 IceSSL . . . 1450
1.12.3.1 Using IceSSL . . . 1452
1.12.3.2 Configuring IceSSL . . . 1455
1.12.3.3 Programming IceSSL . . . 1466
1.12.3.3.1 Programming IceSSL in CPP . . . 1467
1.12.3.3.2 Programming IceSSL in Java . . . 1474

Copyright 2016, ZeroC, Inc.

1.12.3.3.3 Programming IceSSL in .NET . . . 1479
1.12.3.4 Advanced IceSSL Topics . . . 1484
1.12.3.5 Setting up a Certificate Authority . . . 1491
1.13 Ice Services . . . 1492
1.13.1 Freeze . . . 1493
1.13.1.1 Freeze Evictors . . . 1494
1.13.1.1.1 Freeze Evictor Concepts . . . 1495
1.13.1.1.2 Background Save Evictor . . . 1501
1.13.1.1.3 Transactional Evictor . . . 1505
1.13.1.1.4 Using a Freeze Evictor in the File System Server . . . 1510
1.13.1.2 Freeze Maps . . . 1532
1.13.1.2.1 Freeze Map Concepts . . . 1533
1.13.1.2.2 Using a Freeze Map in C++ . . . 1541
1.13.1.2.3 slice2freeze Command-Line Options . . . 1549
1.13.1.2.4 Using a Freeze Map in Java . . . 1551
1.13.1.2.5 slice2freezej Command-Line Options . . . 1563
1.13.1.2.6 Using a Freeze Map in the File System Server . . . 1564
1.13.1.3 Freeze Catalogs . . . 1590
1.13.1.4 Backing Up Freeze Databases . . . 1592
1.13.2 FreezeScript . . . 1593
1.13.2.1 Migrating a Freeze Database . . . 1594
1.13.2.1.1 Automatic Database Migration . . . 1595
1.13.2.1.2 Custom Database Migration . . . 1599
1.13.2.1.3 FreezeScript Transformation XML Reference . . . 1603
1.13.2.1.4 Using transformdb . . . 1609
1.13.2.2 Inspecting a Freeze Database . . . 1615
1.13.2.2.1 Using dumpdb . . . 1616
1.13.2.2.2 FreezeScript Inspection XML Reference . . . 1622
1.13.2.3 FreezeScript Descriptor Expression Language . . . 1628
1.13.3 Glacier2 . . . 1631
1.13.3.1 Common Firewall Traversal Issues . . . 1632
1.13.3.2 About Glacier2 . . . 1633
1.13.3.3 How Glacier2 Works . . . 1634
1.13.3.4 Getting Started with Glacier2 . . . 1635
1.13.3.5 Callbacks through Glacier2 . . . 1642
1.13.3.6 Glacier2 Helper Classes . . . 1645
1.13.3.7 Securing a Glacier2 Router . . . 1652
1.13.3.8 Glacier2 Session Management . . . 1659
1.13.3.9 Dynamic Request Filtering with Glacier2 . . . 1662
1.13.3.10 Glacier2 Request Buffering . . . 1664
1.13.3.11 How Glacier2 uses Request Contexts . . . 1665
1.13.3.12 Configuring Glacier2 behind an External Firewall . . . 1667
1.13.3.13 Advanced Glacier2 Client Configurations . . . 1668
1.13.3.14 IceGrid and Glacier2 Integration . . . 1670
1.13.3.15 Glacier2 Metrics . . . 1672
1.13.4 IceGrid . . . 1674
1.13.4.1 IceGrid Architecture . . . 1676
1.13.4.2 Getting Started with IceGrid . . . 1678
1.13.4.3 Using IceGrid Deployment . . . 1683
1.13.4.4 Well-Known Objects . . . 1691
1.13.4.5 IceGrid Templates . . . 1699
1.13.4.6 IceBox Integration with IceGrid . . . 1705
1.13.4.7 Object Adapter Replication . . . 1709
1.13.4.8 Load Balancing . . . 1712
1.13.4.9 Resource Allocation using IceGrid Sessions . . . 1723
1.13.4.10 Registry Replication . . . 1730
1.13.4.11 Application Distribution . . . 1734
1.13.4.12 IceGrid Administrative Sessions . . . 1741
1.13.4.13 Glacier2 Integration with IceGrid . . . 1748
1.13.4.14 IceGrid XML Reference . . . 1752
1.13.4.14.1 Adapter Descriptor Element . . . 1753
1.13.4.14.2 Allocatable Descriptor Element . . . 1755
1.13.4.14.3 Application Descriptor Element . . . 1756
1.13.4.14.4 DbEnv Descriptor Element . . . 1757
1.13.4.14.5 DbProperty Descriptor Element . . . 1758
1.13.4.14.6 Description Descriptor Element . . . 1759
1.13.4.14.7 Directory Descriptor Element . . . 1760
1.13.4.14.8 Distrib Descriptor Element . . . 1761
1.13.4.14.9 IceBox Descriptor Element . . . 1762
1.13.4.14.10 IceGrid Descriptor Element . . . 1763
1.13.4.14.11 Load-Balancing Descriptor Element . . . 1764
1.13.4.14.12 Log Descriptor Element . . . 1765

Copyright 2016, ZeroC, Inc.

1.13.4.14.13 Node Descriptor Element . . . 1766
1.13.4.14.14 Object Descriptor Element . . . 1767
1.13.4.14.15 Parameter Descriptor Element . . . 1768
1.13.4.14.16 Properties Descriptor Element . . . 1769
1.13.4.14.17 Property Descriptor Element . . . 1770
1.13.4.14.18 Replica-Group Descriptor Element . . . 1771
1.13.4.14.19 Server Descriptor Element . . . 1772
1.13.4.14.20 Server-Instance Descriptor Element . . . 1774
1.13.4.14.21 Server-Template Descriptor Element . . . 1775
1.13.4.14.22 Service Descriptor Element . . . 1776
1.13.4.14.23 Service-Instance Descriptor Element . . . 1777
1.13.4.14.24 Service-Template Descriptor Element . . . 1778
1.13.4.14.25 Variable Descriptor Element . . . 1779
1.13.4.14.26 Using Command Line Options in Descriptors . . . 1780
1.13.4.14.27 Setting Environment Variables in Descriptors . . . 1781
1.13.4.15 Using Descriptor Variables and Parameters . . . 1783
1.13.4.16 IceGrid Property Set Semantics . . . 1789
1.13.4.17 IceGrid XML Features . . . 1794
1.13.4.18 IceGrid Server Reference . . . 1797
1.13.4.18.1 icegridregistry . . . 1798
1.13.4.18.2 icegridnode . . . 1800
1.13.4.18.3 Well-Known Registry Objects . . . 1802
1.13.4.18.4 IceGrid Persistent Data . . . 1804
1.13.4.18.5 Promoting a Registry Slave . . . 1806
1.13.4.19 IceGrid and the Administrative Facility . . . 1807
1.13.4.20 Securing IceGrid . . . 1814
1.13.4.21 icegridadmin Command Line Tool . . . 1818
1.13.4.22 IceGrid Admin Graphical Tool . . . 1824
1.13.4.22.1 Getting Started with IceGrid Admin . . . 1825
1.13.4.22.2 Live Deployment Tab . . . 1828
1.13.4.22.3 Application Tabs . . . 1863
1.13.4.23 IceGrid Server Activation . . . 1886
1.13.4.24 IceGrid Troubleshooting . . . 1889
1.13.4.25 IceGrid Database Utility . . . 1892
1.13.5 IcePatch2 . . . 1893
1.13.5.1 Using icepatch2calc . . . 1894
1.13.5.2 Running the IcePatch2 Server . . . 1896
1.13.5.3 Running the IcePatch2 Client . . . 1897
1.13.5.4 IcePatch2 Object Identities . . . 1899
1.13.5.5 IcePatch2 Client Utility Library . . . 1900
1.13.6 IceStorm . . . 1904
1.13.6.1 IceStorm Concepts . . . 1905
1.13.6.2 IceStorm Interfaces . . . 1907
1.13.6.3 Using IceStorm . . . 1910
1.13.6.3.1 Implementing an IceStorm Publisher . . . 1911
1.13.6.3.2 Using an IceStorm Publisher Object . . . 1915
1.13.6.3.3 Implementing an IceStorm Subscriber . . . 1917
1.13.6.3.4 Publishing to a Specific Subscriber . . . 1924
1.13.6.4 Highly Available IceStorm . . . 1926
1.13.6.5 IceStorm Administration . . . 1930
1.13.6.6 Topic Federation . . . 1932
1.13.6.7 IceStorm Quality of Service . . . 1936
1.13.6.8 IceStorm Delivery Modes . . . 1938
1.13.6.9 Configuring IceStorm . . . 1940
1.13.6.10 IceStorm Metrics . . . 1945
1.13.6.11 IceStorm Database Utility . . . 1947
1.14 The Ice Protocol . . . 1948
1.14.1 Data Encoding . . . 1949
1.14.1.1 Basic Data Encoding . . . 1950
1.14.1.2 Data Encoding for Exceptions . . . 1954
1.14.1.3 Data Encoding for Classes . . . 1959
1.14.1.3.1 Data Encoding for Class Type IDs . . . 1963
1.14.1.3.2 Simple Example of Class Encoding . . . 1965
1.14.1.3.3 Data Encoding for Class Graphs . . . 1971
1.14.1.4 Data Encoding for Interfaces . . . 1980
1.14.1.5 Data Encoding for Proxies . . . 1982
1.14.1.6 Data Encoding for Optional Values . . . 1988
1.14.2 Protocol Messages . . . 1993
1.14.3 Protocol Compression . . . 1999
1.14.4 Protocol and Encoding Versions . . . 2000
1.15 Best Practices . . . 2002
1.15.1 Optional Values . . . 2003

Copyright 2016, ZeroC, Inc.

1.15.2 Server Implementation Techniques . . . 2009
1.15.3 Servant Evictors . . . 2015
1.15.3.1 Implementing a Servant Evictor in C++ . . . 2017
1.15.3.2 Implementing a Servant Evictor in Java . . . 2024
1.15.3.3 Implementing a Servant Evictor in C-Sharp . . . 2032
1.15.4 Object Life Cycle . . . 2041
1.15.4.1 Understanding Object Life Cycle . . . 2042
1.15.4.2 Object Existence and Non-Existence . . . 2043
1.15.4.3 Life Cycle of Proxies, Servants, and Ice Objects . . . 2046
1.15.4.4 Object Creation . . . 2048
1.15.4.5 Object Destruction . . . 2052
1.15.4.5.1 Idempotency and Life Cycle Operations . . . 2054
1.15.4.5.2 Implementing a destroy Operation . . . 2055
1.15.4.5.3 Cleaning Up a Destroyed Servant . . . 2057
1.15.4.5.4 Life Cycle and Collection Operations . . . 2059
1.15.4.5.5 Life Cycle and Normal Operations . . . 2063
1.15.4.6 Removing Cyclic Dependencies . . . 2068
1.15.4.6.1 Acquiring Locks without Deadlocks . . . 2069
1.15.4.6.2 Reaping Objects . . . 2070
1.15.4.7 Object Identity and Uniqueness . . . 2075
1.15.4.8 Object Life Cycle for the File System Application . . . 2077
1.15.4.8.1 Implementing Object Life Cycle in C++ . . . 2079
1.15.4.8.2 Implementing Object Life Cycle in Java . . . 2087
1.15.4.9 Avoiding Server-Side Garbage . . . 2097
1.16 Platform-Specific Features . . . 2104
1.16.1 Windows Services . . . 2105
1.16.1.1 Installing a Windows Service . . . 2106
1.16.1.2 Using the Ice Service Installer . . . 2107
1.16.1.3 Manually Installing a Service as a Windows Service . . . 2110
1.16.1.4 Troubleshooting Windows Services . . . 2116
1.16.2 Windows Store Applications . . . 2118
1.17 Property Reference . . . 2120
1.17.1 Object Adapter Properties . . . 2121
1.17.2 Proxy Properties . . . 2126
1.17.3 Miscellaneous Ice.* Properties . . . 2129
1.17.4 Freeze.* . . . 2137
1.17.5 Glacier2.* . . . 2144
1.17.6 Ice.ACM.* . . . 2154
1.17.7 Ice.Admin.* . . . 2158
1.17.8 Ice.Config . . . 2161
1.17.9 Ice.Default.* . . . 2162
1.17.10 Ice.InitPlugins . . . 2166
1.17.11 Ice.IPv4 . . . 2167
1.17.12 Ice.IPv6 . . . 2168
1.17.13 Ice.Override.* . . . 2169
1.17.14 Ice.Plugin.* . . . 2171
1.17.15 Ice.PluginLoadOrder . . . 2175
1.17.16 Ice.PreferIPv6Address . . . 2176
1.17.17 Ice.TCP.* . . . 2177
1.17.18 Ice.ThreadPool.* . . . 2178
1.17.19 Ice.ThreadPriority . . . 2181
1.17.20 Ice.Trace.* . . . 2182
1.17.21 Ice.UDP.* . . . 2185
1.17.22 Ice.Warn.* . . . 2186
1.17.23 IceBox.* . . . 2188
1.17.24 IceBoxAdmin . . . 2192
1.17.25 IceDiscovery.* . . . 2193
1.17.26 IceGrid.* . . . 2196
1.17.27 IceGridAdmin.* . . . 2210
1.17.28 IceLocatorDiscovery.* . . . 2213
1.17.29 IceMX.Metrics.* . . . 2216
1.17.30 IcePatch2.* . . . 2218
1.17.31 IcePatch2Client.* . . . 2219
1.17.32 IceSSL.* . . . 2221
1.17.33 IceStorm Properties . . . 2238
1.17.34 IceStormAdmin.* . . . 2244
2. Ice Release Notes . . . 2246
2.1 Supported Platforms for Ice 3.6.3 . . . 2247
2.2 New Features in Ice 3.6 . . . 2250
2.3 Backward Compatibility of Ice Versions . . . 2259
2.4 Upgrading your Application from Ice 3.6.0 or later . . . 2261
2.5 Upgrading your Application from Ice 3.5 . . . 2264

Copyright 2016, ZeroC, Inc.

2.6 Upgrading your Application from Ice 3.4 . . . 2279
2.7 Upgrading your Application from Ice 3.3 . . . 2285
2.8 Upgrading your Application from Ice 3.2 or Earlier Releases . . . 2302
2.9 Platform Notes for Ice 3.6 . . . 2308
2.10 Using the Windows Binary Distribution . . . 2310
2.11 Using the Linux Binary Distributions . . . 2318
2.12 Using the macOS Binary Distribution . . . 2329
2.13 Using Ice on Android . . . 2333
2.14 Using the JavaScript Distribution . . . 2335
2.15 Using the Python Distribution . . . 2338
2.16 Using the Ruby Distribution . . . 2340
2.17 Getting Started with Ice on AWS . . . 2342
2.18 Building Third Party Packages from Source . . . 2344
2.18.1 Building Berkeley DB, bzip2 and mcpp on Windows . . . 2345
2.18.2 Building Berkeley DB and mcpp on Linux . . . 2347
2.18.3 Building Berkeley DB on macOS . . . 2349
3. Ice-E Release Notes . . . 2350
3.1 Supported Platforms for Ice-E 3.6.3 . . . 2351
3.2 New Features in Ice-E 3.6.3 . . . 2352
3.3 Ice-E Feature Set . . . 2353
3.4 Ice-E Application Notes . . . 2354
3.5 Upgrading your Ice-E Application . . . 2355
3.6 Using the Ice-E Debian Binary Distribution . . . 2356
3.7 Using Ice-E with Yocto . . . 2360
4. Ice Touch Release Notes . . . 2362
4.1 Supported Platforms for Ice Touch 3.6.3 . . . 2363
4.2 New Features in Ice Touch 3.6.3 . . . 2364
4.3 Ice Touch Feature Set . . . 2365
4.4 Upgrading your Ice Touch Application . . . 2367
4.5 Using the Ice Touch Binary Distribution . . . 2368
5. Slice API Reference . . . 2370
5.1 Freeze Slice API . . . 2371
5.1.1 Freeze-BackgroundSaveEvictor . . . 2372
5.1.2 Freeze-CatalogData . . . 2374
5.1.3 Freeze-Connection . . . 2375
5.1.4 Freeze-DatabaseException . . . 2377
5.1.5 Freeze-DeadlockException . . . 2378
5.1.6 Freeze-Evictor . . . 2379
5.1.7 Freeze-EvictorDeactivatedException . . . 2383
5.1.8 Freeze-EvictorIterator . . . 2384
5.1.9 Freeze-IndexNotFoundException . . . 2385
5.1.10 Freeze-InvalidPositionException . . . 2386
5.1.11 Freeze-NoSuchElementException . . . 2387
5.1.12 Freeze-NotFoundException . . . 2388
5.1.13 Freeze-ObjectRecord . . . 2389
5.1.14 Freeze-ServantInitializer . . . 2390
5.1.15 Freeze-Statistics . . . 2391
5.1.16 Freeze-Transaction . . . 2392
5.1.17 Freeze-TransactionalEvictor . . . 2393
5.1.18 Freeze-TransactionAlreadyInProgressException . . . 2394
5.2 Glacier2 Slice API . . . 2395
5.2.1 Glacier2-CannotCreateSessionException . . . 2396
5.2.2 Glacier2-IdentitySet . . . 2397
5.2.3 Glacier2-PermissionDeniedException . . . 2398
5.2.4 Glacier2-PermissionsVerifier . . . 2399
5.2.5 Glacier2-Router . . . 2400
5.2.6 Glacier2-Session . . . 2402
5.2.7 Glacier2-SessionControl . . . 2403
5.2.8 Glacier2-SessionManager . . . 2404
5.2.9 Glacier2-SessionNotExistException . . . 2405
5.2.10 Glacier2-SSLInfo . . . 2406
5.2.11 Glacier2-SSLPermissionsVerifier . . . 2407
5.2.12 Glacier2-SSLSessionManager . . . 2408
5.2.13 Glacier2-StringSet . . . 2409
5.3 IceBox Slice API . . . 2410
5.3.1 IceBox-AlreadyStartedException . . . 2411
5.3.2 IceBox-AlreadyStoppedException . . . 2412
5.3.3 IceBox-FailureException . . . 2413
5.3.4 IceBox-NoSuchServiceException . . . 2414
5.3.5 IceBox-Service . . . 2415
5.3.6 IceBox-ServiceManager . . . 2416
5.3.7 IceBox-ServiceObserver . . . 2417

Copyright 2016, ZeroC, Inc.

5.4 IceGrid Slice API . . . 2418
5.4.1 IceGrid-AccessDeniedException . . . 2424
5.4.2 IceGrid-AdapterDescriptor . . . 2425
5.4.3 IceGrid-AdapterDynamicInfo . . . 2427
5.4.4 IceGrid-AdapterInfo . . . 2428
5.4.5 IceGrid-AdapterNotExistException . . . 2429
5.4.6 IceGrid-AdapterObserver . . . 2430
5.4.7 IceGrid-AdaptiveLoadBalancingPolicy . . . 2431
5.4.8 IceGrid-Admin . . . 2432
5.4.9 IceGrid-AdminSession . . . 2444
5.4.10 IceGrid-AllocationException . . . 2449
5.4.11 IceGrid-AllocationTimeoutException . . . 2450
5.4.12 IceGrid-ApplicationDescriptor . . . 2451
5.4.13 IceGrid-ApplicationInfo . . . 2453
5.4.14 IceGrid-ApplicationNotExistException . . . 2454
5.4.15 IceGrid-ApplicationObserver . . . 2455
5.4.16 IceGrid-ApplicationUpdateDescriptor . . . 2456
5.4.17 IceGrid-ApplicationUpdateInfo . . . 2458
5.4.18 IceGrid-BadSignalException . . . 2459
5.4.19 IceGrid-BoxedDistributionDescriptor . . . 2460
5.4.20 IceGrid-BoxedString . . . 2461
5.4.21 IceGrid-CommunicatorDescriptor . . . 2462
5.4.22 IceGrid-DbEnvDescriptor . . . 2463
5.4.23 IceGrid-DeploymentException . . . 2464
5.4.24 IceGrid-DistributionDescriptor . . . 2465
5.4.25 IceGrid-FileIterator . . . 2466
5.4.26 IceGrid-FileNotAvailableException . . . 2467
5.4.27 IceGrid-FileParser . . . 2468
5.4.28 IceGrid-IceBoxDescriptor . . . 2469
5.4.29 IceGrid-LoadBalancingPolicy . . . 2470
5.4.30 IceGrid-LoadInfo . . . 2471
5.4.31 IceGrid-LoadSample . . . 2472
5.4.32 IceGrid-Locator . . . 2473
5.4.33 IceGrid-NodeDescriptor . . . 2474
5.4.34 IceGrid-NodeDynamicInfo . . . 2475
5.4.35 IceGrid-NodeInfo . . . 2476
5.4.36 IceGrid-NodeNotExistException . . . 2477
5.4.37 IceGrid-NodeObserver . . . 2478
5.4.38 IceGrid-NodeUnreachableException . . . 2480
5.4.39 IceGrid-NodeUpdateDescriptor . . . 2481
5.4.40 IceGrid-ObjectDescriptor . . . 2483
5.4.41 IceGrid-ObjectExistsException . . . 2484
5.4.42 IceGrid-ObjectInfo . . . 2485
5.4.43 IceGrid-ObjectNotRegisteredException . . . 2486
5.4.44 IceGrid-ObjectObserver . . . 2487
5.4.45 IceGrid-ObserverAlreadyRegisteredException . . . 2488
5.4.46 IceGrid-OrderedLoadBalancingPolicy . . . 2489
5.4.47 IceGrid-ParseException . . . 2490
5.4.48 IceGrid-PatchException . . . 2491
5.4.49 IceGrid-PermissionDeniedException . . . 2492
5.4.50 IceGrid-PropertyDescriptor . . . 2493
5.4.51 IceGrid-PropertySetDescriptor . . . 2494
5.4.52 IceGrid-Query . . . 2495
5.4.53 IceGrid-RandomLoadBalancingPolicy . . . 2497
5.4.54 IceGrid-Registry . . . 2498
5.4.55 IceGrid-RegistryInfo . . . 2501
5.4.56 IceGrid-RegistryNotExistException . . . 2502
5.4.57 IceGrid-RegistryObserver . . . 2503
5.4.58 IceGrid-RegistryPluginFacade . . . 2504
5.4.59 IceGrid-RegistryUnreachableException . . . 2508
5.4.60 IceGrid-ReplicaGroupDescriptor . . . 2509
5.4.61 IceGrid-ReplicaGroupFilter . . . 2510
5.4.62 IceGrid-RoundRobinLoadBalancingPolicy . . . 2511
5.4.63 IceGrid-ServerDescriptor . . . 2512
5.4.64 IceGrid-ServerDynamicInfo . . . 2514
5.4.65 IceGrid-ServerInfo . . . 2515
5.4.66 IceGrid-ServerInstanceDescriptor . . . 2516
5.4.67 IceGrid-ServerNotExistException . . . 2517
5.4.68 IceGrid-ServerStartException . . . 2518
5.4.69 IceGrid-ServerState . . . 2519
5.4.70 IceGrid-ServerStopException . . . 2520
5.4.71 IceGrid-ServerUnreachableException . . . 2521

Copyright 2016, ZeroC, Inc.

5.4.72 IceGrid-ServiceDescriptor . . . 2522
5.4.73 IceGrid-ServiceInstanceDescriptor . . . 2523
5.4.74 IceGrid-Session . . . 2524
5.4.75 IceGrid-TemplateDescriptor . . . 2526
5.4.76 IceGrid-TypeFilter . . . 2527
5.4.77 IceGrid-UserAccountMapper . . . 2528
5.4.78 IceGrid-UserAccountNotFoundException . . . 2529
5.5 IceLocatorDiscovery Slice API . . . 2530
5.5.1 IceLocatorDiscovery-Lookup . . . 2531
5.5.2 IceLocatorDiscovery-LookupReply . . . 2532
5.6 IceMX Slice API . . . 2533
5.6.1 IceMX-ChildInvocationMetrics . . . 2535
5.6.2 IceMX-CollocatedMetrics . . . 2536
5.6.3 IceMX-ConnectionMetrics . . . 2537
5.6.4 IceMX-DispatchMetrics . . . 2538
5.6.5 IceMX-InvocationMetrics . . . 2539
5.6.6 IceMX-Metrics . . . 2540
5.6.7 IceMX-MetricsAdmin . . . 2541
5.6.8 IceMX-MetricsFailures . . . 2543
5.6.9 IceMX-RemoteMetrics . . . 2544
5.6.10 IceMX-SessionMetrics . . . 2545
5.6.11 IceMX-SubscriberMetrics . . . 2546
5.6.12 IceMX-ThreadMetrics . . . 2547
5.6.13 IceMX-TopicMetrics . . . 2548
5.6.14 IceMX-UnknownMetricsView . . . 2549
5.7 IcePatch2 Slice API . . . 2550
5.7.1 IcePatch2-FileAccessException . . . 2552
5.7.2 IcePatch2-FileInfo . . . 2553
5.7.3 IcePatch2-FileServer . . . 2554
5.7.4 IcePatch2-FileSizeRangeException . . . 2556
5.7.5 IcePatch2-LargeFileInfo . . . 2557
5.7.6 IcePatch2-PartitionOutOfRangeException . . . 2558
5.8 Ice Slice API . . . 2559
5.8.1 Ice-ACM . . . 2566
5.8.2 Ice-ACMClose . . . 2567
5.8.3 Ice-ACMHeartbeat . . . 2568
5.8.4 Ice-AdapterAlreadyActiveException . . . 2569
5.8.5 Ice-AdapterNotFoundException . . . 2570
5.8.6 Ice-AlreadyRegisteredException . . . 2571
5.8.7 Ice-BadMagicException . . . 2572
5.8.8 Ice-CloneNotImplementedException . . . 2573
5.8.9 Ice-CloseConnectionException . . . 2574
5.8.10 Ice-CloseTimeoutException . . . 2575
5.8.11 Ice-CollocationOptimizationException . . . 2576
5.8.12 Ice-Communicator . . . 2577
5.8.13 Ice-CommunicatorDestroyedException . . . 2585
5.8.14 Ice-CompressionException . . . 2586
5.8.15 Ice-ConnectFailedException . . . 2587
5.8.16 Ice-Connection . . . 2588
5.8.17 Ice-ConnectionCallback . . . 2591
5.8.18 Ice-ConnectionInfo . . . 2592
5.8.19 Ice-ConnectionLostException . . . 2593
5.8.20 Ice-ConnectionNotValidatedException . . . 2594
5.8.21 Ice-ConnectionRefusedException . . . 2595
5.8.22 Ice-ConnectionTimeoutException . . . 2596
5.8.23 Ice-ConnectTimeoutException . . . 2597
5.8.24 Ice-Current . . . 2598
5.8.25 Ice-DatagramLimitException . . . 2600
5.8.26 Ice-DNSException . . . 2601
5.8.27 Ice-EncapsulationException . . . 2602
5.8.28 Ice-EncodingVersion . . . 2603
5.8.29 Ice-Endpoint . . . 2604
5.8.30 Ice-EndpointInfo . . . 2605
5.8.31 Ice-EndpointParseException . . . 2606
5.8.32 Ice-EndpointSelectionType . . . 2607
5.8.33 Ice-EndpointSelectionTypeParseException . . . 2608
5.8.34 Ice-FacetNotExistException . . . 2609
5.8.35 Ice-FeatureNotSupportedException . . . 2610
5.8.36 Ice-FileException . . . 2611
5.8.37 Ice-FixedProxyException . . . 2612
5.8.38 Ice-ForcedCloseConnectionException . . . 2613
5.8.39 Ice-Identity . . . 2614

Copyright 2016, ZeroC, Inc.

5.8.40 Ice-IdentityParseException . . . 2616
5.8.41 Ice-IllegalIdentityException . . . 2617
5.8.42 Ice-IllegalMessageSizeException . . . 2618
5.8.43 Ice-IllegalServantException . . . 2619
5.8.44 Ice-ImplicitContext . . . 2620
5.8.45 Ice-InitializationException . . . 2622
5.8.46 Ice-Instrumentation . . . 2623
5.8.46.1 Ice-Instrumentation-ChildInvocationObserver . . . 2624
5.8.46.2 Ice-Instrumentation-CollocatedObserver . . . 2625
5.8.46.3 Ice-Instrumentation-CommunicatorObserver . . . 2626
5.8.46.4 Ice-Instrumentation-ConnectionObserver . . . 2628
5.8.46.5 Ice-Instrumentation-ConnectionState . . . 2629
5.8.46.6 Ice-Instrumentation-DispatchObserver . . . 2630
5.8.46.7 Ice-Instrumentation-InvocationObserver . . . 2631
5.8.46.8 Ice-Instrumentation-Observer . . . 2632
5.8.46.9 Ice-Instrumentation-ObserverUpdater . . . 2633
5.8.46.10 Ice-Instrumentation-RemoteObserver . . . 2634
5.8.46.11 Ice-Instrumentation-ThreadObserver . . . 2635
5.8.46.12 Ice-Instrumentation-ThreadState . . . 2636
5.8.47 Ice-InvalidReplicaGroupIdException . . . 2637
5.8.48 Ice-InvocationCanceledException . . . 2638
5.8.49 Ice-InvocationTimeoutException . . . 2639
5.8.50 Ice-IPConnectionInfo . . . 2640
5.8.51 Ice-IPEndpointInfo . . . 2641
5.8.52 Ice-Locator . . . 2642
5.8.53 Ice-LocatorFinder . . . 2644
5.8.54 Ice-LocatorRegistry . . . 2645
5.8.55 Ice-Logger . . . 2647
5.8.56 Ice-LoggerAdmin . . . 2649
5.8.57 Ice-LogMessage . . . 2651
5.8.58 Ice-LogMessageType . . . 2652
5.8.59 Ice-MarshalException . . . 2653
5.8.60 Ice-MemoryLimitException . . . 2654
5.8.61 Ice-NoEndpointException . . . 2655
5.8.62 Ice-NoObjectFactoryException . . . 2656
5.8.63 Ice-NotRegisteredException . . . 2657
5.8.64 Ice-ObjectAdapter . . . 2658
5.8.65 Ice-ObjectAdapterDeactivatedException . . . 2668
5.8.66 Ice-ObjectAdapterIdInUseException . . . 2669
5.8.67 Ice-ObjectFactory . . . 2670
5.8.68 Ice-ObjectNotExistException . . . 2671
5.8.69 Ice-ObjectNotFoundException . . . 2672
5.8.70 Ice-OpaqueEndpointInfo . . . 2673
5.8.71 Ice-OperationInterruptedException . . . 2674
5.8.72 Ice-OperationMode . . . 2675
5.8.73 Ice-OperationNotExistException . . . 2676
5.8.74 Ice-Plugin . . . 2677
5.8.75 Ice-PluginInitializationException . . . 2678
5.8.76 Ice-PluginManager . . . 2679
5.8.77 Ice-Process . . . 2681
5.8.78 Ice-Properties . . . 2682
5.8.79 Ice-PropertiesAdmin . . . 2686
5.8.80 Ice-ProtocolException . . . 2687
5.8.81 Ice-ProtocolVersion . . . 2688
5.8.82 Ice-ProxyParseException . . . 2689
5.8.83 Ice-ProxyUnmarshalException . . . 2690
5.8.84 Ice-RemoteLogger . . . 2691
5.8.85 Ice-RemoteLoggerAlreadyAttachedException . . . 2692
5.8.86 Ice-RequestFailedException . . . 2693
5.8.87 Ice-ResponseSentException . . . 2694
5.8.88 Ice-Router . . . 2695
5.8.89 Ice-RouterFinder . . . 2696
5.8.90 Ice-SecurityException . . . 2697
5.8.91 Ice-ServantLocator . . . 2698
5.8.92 Ice-ServerNotFoundException . . . 2700
5.8.93 Ice-SocketException . . . 2701
5.8.94 Ice-StringConversionException . . . 2702
5.8.95 Ice-SyscallException . . . 2703
5.8.96 Ice-TCPConnectionInfo . . . 2704
5.8.97 Ice-TCPEndpointInfo . . . 2705
5.8.98 Ice-TimeoutException . . . 2706
5.8.99 Ice-TwowayOnlyException . . . 2707

Copyright 2016, ZeroC, Inc.

5.8.100 Ice-UDPConnectionInfo . . . 2708
5.8.101 Ice-UDPEndpointInfo . . . 2709
5.8.102 Ice-UnexpectedObjectException . . . 2710
5.8.103 Ice-UnknownException . . . 2711
5.8.104 Ice-UnknownLocalException . . . 2712
5.8.105 Ice-UnknownMessageException . . . 2713
5.8.106 Ice-UnknownReplyStatusException . . . 2714
5.8.107 Ice-UnknownRequestIdException . . . 2715
5.8.108 Ice-UnknownUserException . . . 2716
5.8.109 Ice-UnmarshalOutOfBoundsException . . . 2717
5.8.110 Ice-UnsupportedEncodingException . . . 2718
5.8.111 Ice-UnsupportedProtocolException . . . 2719
5.8.112 Ice-VersionMismatchException . . . 2720
5.8.113 Ice-VersionParseException . . . 2721
5.8.114 Ice-WSConnectionInfo . . . 2722
5.8.115 Ice-WSEndpointInfo . . . 2723
5.9 IceSSL Slice API . . . 2724
5.9.1 IceSSL-ConnectionInfo . . . 2725
5.9.2 IceSSL-EndpointInfo . . . 2726
5.9.3 IceSSL-WSSConnectionInfo . . . 2727
5.9.4 IceSSL-WSSEndpointInfo . . . 2728
5.10 IceStorm Slice API . . . 2729
5.10.1 IceStorm-AlreadySubscribed . . . 2731
5.10.2 IceStorm-BadQoS . . . 2732
5.10.3 IceStorm-Finder . . . 2733
5.10.4 IceStorm-InvalidSubscriber . . . 2734
5.10.5 IceStorm-LinkExists . . . 2735
5.10.6 IceStorm-LinkInfo . . . 2736
5.10.7 IceStorm-NoSuchLink . . . 2737
5.10.8 IceStorm-NoSuchTopic . . . 2738
5.10.9 IceStorm-Topic . . . 2739
5.10.10 IceStorm-TopicExists . . . 2742
5.10.11 IceStorm-TopicManager . . . 2743

Copyright 2016, ZeroC, Inc.

Ice Manual

Distributed Programming with Ice

The Internet Communications Engine (Ice) is a modern object-oriented toolkit that enables you to build distributed applications with minimal effort. Ice allows you to focus your efforts on your application logic, and it takes care of all interactions with low-level network programming interfaces. With Ice, there is no need to worry about details such as opening network connections, serializing and deserializing data for network transmission, or retrying failed connection attempts.

The main design goals of Ice are:

- Provide an object-oriented middleware platform suitable for use in heterogeneous environments.
- Provide a full set of features that support development of realistic distributed applications for a wide variety of domains.
- Avoid unnecessary complexity, making the platform easy to learn and to use.
- Provide an implementation that is efficient in network bandwidth, memory use, and CPU overhead.
- Provide an implementation that has built-in security, making it suitable for use over insecure public networks.

In simpler terms, the Ice design goals could be stated as "Let's build a powerful middleware platform that makes the developer's life easier."

The acronym "Ice" is pronounced as a single syllable, like the word for frozen water.

Getting Help with Ice

If you have a question and you cannot find an answer in this manual, you can visit our developer forums to see if another developer has encountered the same issue. If you still need help, feel free to post your question on the forum, which ZeroC's developers monitor regularly. Note, however, that we can provide only limited free support in our forums. For guaranteed response and problem resolution times, we highly recommend purchasing commercial support.

Feedback about the Manual

We would very much like to hear from you in case you find any bugs (however minor) in this manual. We also would like to hear your opinion on the contents, and any suggestions as to how it might be improved. You can contact us via e-mail at icebook@zeroc.com.

Legal Notices

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book and ZeroC was aware of the trademark claim, the designations have been printed in initial caps or all caps. The authors and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

License

This manual is provided under one of two licenses, whichever you prefer:

- Creative Commons Attribution-No Derivative Works 3.0 Unported License.
  This license does not permit you to make modifications.
- Creative Commons Attribution-Noncommercial-Share Alike 3.0 Unported License.
  This license permits you to make modifications, but for non-commercial use only. If you distribute this manual under this license, you must prominently include the following text:

  > *This document is derived from ZeroC's* Ice Manual*, Copyright © ZeroC, Inc. 2003-2015.*

Copyright 2016, ZeroC, Inc.

> *You can find the latest version of this document at:*
> *https://doc.zeroc.com/display/Ice/Ice+Manual*

Copyright

Copyright © 2003-2016 by ZeroC, Inc.
mailto:info@zeroc.com
https://zeroc.com

Ice Overview

The following topics provide a high-level overview of Ice:

- Ice Architecture introduces fundamental concepts and terminology, and outlines how Slice definitions, language mappings, and the Ice run time and protocol work in concert to create clients and servers.
- Ice Services Overview briefly presents the object services provided by Ice.
- Architectural Benefits of Ice outlines the benefits that result from the Ice architecture.

Topics

- Ice Architecture
- Ice Services Overview
- Architectural Benefits of Ice

Copyright 2016, ZeroC, Inc.

Ice Architecture

Ice is an object-oriented middleware platform. Fundamentally, this means that Ice provides tools, APIs, and library support for building object-oriented client-server applications. Ice applications are suitable for use in heterogeneous environments: client and server can be written in different programming languages, can run on different operating systems and machine architectures, and can communicate using a variety of networking technologies. The source code for these applications is portable regardless of the deployment environment.

Topics:

- Terminology
- Slice (Specification Language for Ice)
- Overview of the Language Mappings
- Client and Server Structure
- Overview of the Ice Protocol

See Also

- Ice Services Overview
- Architectural Benefits of Ice

Copyright 2016, ZeroC, Inc.

Terminology

Every computing technology creates its own vocabulary as it evolves. Ice is no exception. However, the amount of new jargon used by Ice is minimal. Rather than inventing new terms, we have used existing terminology as much as possible. If you have used another middleware technology in the past, you will be familiar with much of what follows. (However, we suggest you at least skim the material because a few terms used by Ice *do* differ from the corresponding terms used by other middleware.)

On this page:

- Clients and Servers
- Ice Objects
- Proxies
- Stringified Proxies
- Direct Proxies
- Indirect Proxies
- Direct Versus Indirect Binding
- Fixed Proxies
- Routed Proxies
- Replication
- Replica Groups
- Servants
- At-Most-Once Semantics
- Synchronous Method Invocation
- Asynchronous Method Invocation
- Asynchronous Method Dispatch
- Oneway Method Invocation
- Batched Oneway Method Invocation
- Datagram Invocations
- Batched Datagram Invocations
- Run-Time Exceptions
- User Exceptions
- Properties

Clients and Servers

The terms *client* and *server* are not firm designations for particular parts of an application; rather, they denote roles that are taken by parts of an application for the duration of a request:

- Clients are active entities. They issue requests for service to servers.
- Servers are passive entities. They provide services in response to client requests.

Frequently, servers are not "pure" servers, in the sense that they never issue requests and only respond to requests. Instead, servers often act as a server on behalf of some client but, in turn, act as a client to another server in order to satisfy their client's request.

Similarly, clients often are not "pure" clients, in the sense that they only request service from an object. Instead, clients are frequently client-server hybrids. For example, a client might start a long-running operation on a server; as part of starting the operation, the client can provide a *callback object* to the server that is used by the server to notify the client when the operation is complete. In that case, the client acts as a client when it starts the operation, and as a server when it is notified that the operation is complete.

Such role reversal is common in many systems, so, frequently, client-server systems could be more accurately described as *peer-to-peer* sy stems.

Ice Objects

An *Ice object* is a conceptual entity, or abstraction. An Ice object can be characterized by the following points:

- An Ice object is an entity in the local or a remote address space that can respond to client requests.
- A single Ice object can be instantiated in a single server or, redundantly, in multiple servers. If an object has multiple simultaneous instantiations, it is still a single Ice object.
- Each Ice object has one or more *interfaces*. An interface is a collection of named *operations* that are supported by an object. Clients issue requests by invoking operations.
- An operation has zero or more *parameters* as well as a *return value*. Parameters and return values have a specific *type*. Parameters are named and have a direction: in-parameters are initialized by the client and passed to the server; out-parameters are initialized by the server and passed to the client. (The return value is simply a special out-parameter.)
- An Ice object has a distinguished interface, known as its *main interface*. In addition, an Ice object can provide zero or more alternate interfaces, known as *facets*. Clients can select among the facets of an object to choose the interface they want to work with.
- Each Ice object has a unique *object identity*. An object's identity is an identifying value that distinguishes the object from all other objects. The Ice object model assumes that object identities are globally unique, that is, no two objects within an Ice communication

Copyright 2016, ZeroC, Inc.

domain can have the same object identity.

In practice, you need not use object identities that are globally unique, such as UUIDs, only identities that do not clash with any other identity within your domain of interest. However, there are architectural advantages to using globally unique identifiers, which we explore in our discussion of object life cycle.

Proxies

For a client to be able to contact an Ice object, the client must hold a *proxy* for the Ice object. A proxy is an artifact that is local to the client's address space; it represents the (possibly remote) Ice object for the client. A proxy acts as the local ambassador for an Ice object: when the client invokes an operation on the proxy, the Ice run time:

1. Locates the Ice object
2. Activates the Ice object's server if it is not running
3. Activates the Ice object within the server
4. Transmits any in-parameters to the Ice object
5. Waits for the operation to complete
6. Returns any out-parameters and the return value to the client (or throws an exception in case of an error)

A proxy encapsulates all the necessary information for this sequence of steps to take place. In particular, a proxy contains:

- Addressing information that allows the client-side run time to contact the correct server
- An object identity that identifies which particular object in the server is the target of a request
- An optional facet identifier that determines which particular facet of an object the proxy refers to

Stringified Proxies

The information in a proxy can be expressed as a string. For example, the string:

```
SimplePrinter:default -p 10000
```

is a human-readable representation of a proxy. The Ice run time provides API calls that allow you to convert a proxy to its stringified form and vice versa. This is useful, for example, to store proxies in database tables or text files.

Provided that a client knows the identity of an Ice object and its addressing information, it can create a proxy "out of thin air" by supplying that information. In other words, no part of the information inside a proxy is considered opaque; a client needs to know only an object's identity, addressing information, and (to be able to invoke an operation) the object's type in order to contact the object.

Direct Proxies

A *direct proxy* is a proxy that embeds an object's identity, together with the address at which its server runs. The address is completely specified by:

- a protocol identifier (such TCP/IP or UDP)
- a protocol-specific address (such as a host name and port number)

To contact the object denoted by a direct proxy, the Ice run time uses the addressing information in the proxy to contact the server; the identity of the object is sent to the server with each request made by the client.

Indirect Proxies

An *indirect proxy* has two forms. It may provide only an object's identity, or it may specify an identity together with an object adapter identifier. An object that is accessible using only its identity is called a well-known object. For example, the string:

```
SimplePrinter
```

is a valid proxy for a well-known object with the identity `SimplePrinter`.

An indirect proxy that includes an object adapter identifier has the stringified form

Copyright 2016, ZeroC, Inc.

```
SimplePrinter@PrinterAdapter
```

Any object of the object adapter can be accessed using such a proxy, regardless of whether that object is also a well-known object.

Notice that an indirect proxy contains no addressing information. To determine the correct server, the client-side run time passes the proxy information to a location service. In turn, the location service uses the object identity or the object adapter identifier as the key in a lookup table that contains the address of the server and returns the current server address to the client. The client-side run time now knows how to contact the server and dispatches the client request as usual.

The entire process is similar to the mapping from Internet domain names to IP address by the Domain Name Service (DNS): when we use a domain name, such as `www.zeroc.com`, to look up a web page, the host name is first resolved to an IP address behind the scenes and, once the correct IP address is known, the IP address is used to connect to the server. With Ice, the mapping is from an object identity or object adapter identifier to a protocol-address pair, but otherwise very similar. The client-side run time knows how to contact the location service via configuration (just as web browsers know which DNS server to use via configuration).

Direct Versus Indirect Binding

The process of resolving the information in a proxy to protocol-address pair is known as *binding*. Not surprisingly, *direct binding* is used for direct proxies, and *indirect binding* is used for indirect proxies.

The main advantage of indirect binding is that it allows us to move servers around (that is, change their address) without invalidating existing proxies that are held by clients. In other words, direct proxies avoid the extra lookup to locate the server but no longer work if a server is moved to a different machine. On the other hand, indirect proxies continue to work even if we move (or *migrate*) a server.

Fixed Proxies

A *fixed proxy* is a proxy that is bound to a particular connection: instead of containing addressing information or an adapter name, the proxy contains a connection handle. The connection handle stays valid only for as long as the connection stays open so, once the connection is closed, the proxy no longer works (and will never work again). Fixed proxies cannot be marshaled, that is, they cannot be passed as parameters on operation invocations. Fixed proxies are used to allow bidirectional communication, so a server can make callbacks to a client without having to open a new connection.

Routed Proxies

A *routed proxy* is a proxy that forwards all invocations to a specific target object, instead of sending invocations directly to the actual target. Routed proxies are useful for implementing services such as Glacier2, which enables clients to communicate with servers that are behind a firewall.

Replication

In Ice, *replication* involves making object adapters (and their objects) available at multiple addresses. The goal of replication is usually to provide redundancy by running the same server on several computers. If one of the computers should happen to fail, a server still remains available on the others.

The use of replication implies that applications are designed for it. In particular, it means a client can access an object via one address and obtain the same result as from any other address. Either these objects are stateless, or their implementations are designed to synchronize with a database (or each other) in order to maintain a consistent view of each object's state.

Ice supports a limited form of replication when a proxy specifies multiple addresses for an object. The Ice run time selects one of the addresses at random for its initial connection attempt and tries all of them in the case of a failure. For example, consider this proxy:

```
SimplePrinter:tcp -h server1 -p 10001:tcp -h server2 -p 10002
```

The proxy states that the object with identity `SimplePrinter` is available using TCP at two addresses, one on the host `server1` and another on the host `server2`. The burden falls to users or system administrators to ensure that the servers are actually running on these computers at the specified ports.

Copyright 2016, ZeroC, Inc.

Replica Groups

In addition to the proxy-based replication described above, Ice supports a more useful form of replication known as *replica groups* that requires the use of a location service.

A replica group has a unique identifier and consists of any number of object adapters. An object adapter may be a member of at most one replica group; such an adapter is considered to be a *replicated object adapter*.

After a replica group has been established, its identifier can be used in an indirect proxy in place of an adapter identifier. For example, a replica group identified as `PrinterAdapters` can be used in a proxy as shown below:

```
SimplePrinter@PrinterAdapters
```

The replica group is treated by the location service as a "virtual object adapter." The behavior of the location service when resolving an indirect proxy containing a replica group id is an implementation detail. For example, the location service could decide to return the addresses of all object adapters in the group, in which case the client's Ice run time would select one of the addresses at random using the limited form of replication discussed earlier. Another possibility is for the location service to return only one address, which it decided upon using some heuristic.

Regardless of the way in which a location service resolves a replica group, the key benefit is indirection: the location service as a middleman can add more intelligence to the binding process.

Servants

As we mentioned, an Ice Object is a conceptual entity that has a type, identity, and addressing information. However, client requests ultimately must end up with a concrete server-side processing entity that can provide the behavior for an operation invocation. To put this differently, a client request must ultimately end up executing code inside the server, with that code written in a specific programming language and executing on a specific processor.

The server-side artifact that provides behavior for operation invocations is known as a *servant*. A servant provides substance for (or *incarnates*) one or more Ice objects. In practice, a servant is simply an instance of a class that is written by the server developer and that is registered with the server-side run time as the servant for one or more Ice objects. Methods on the class correspond to the operations on the Ice object's interface and provide the behavior for the operations.

A single servant can incarnate a single Ice object at a time or several Ice objects simultaneously. If the former, the identity of the Ice object incarnated by the servant is implicit in the servant. If the latter, the servant is provided the identity of the Ice object with each request, so it can decide which object to incarnate for the duration of the request.

Conversely, a single Ice object can have multiple servants. For example, we might choose to create a proxy for an Ice object with two different addresses for different machines. In that case, we will have two servers, with each server containing a servant for the same Ice object. When a client invokes an operation on such an Ice object, the client-side run time sends the request to exactly one server. In other words, multiple servants for a single Ice object allow you to build redundant systems: the client-side run time attempts to send the request to one server and, if that attempt fails, sends the request to the second server. An error is reported back to the client-side application code only if that second attempt also fails.

At-Most-Once Semantics

Ice requests have *at-most-once* semantics: the Ice run time does its best to deliver a request to the correct destination and, depending on the exact circumstances, may retry a failed request. Ice guarantees that it will either deliver the request, or, if it cannot deliver the request, inform the client with an appropriate exception; under no circumstances is a request delivered twice, that is, retries are attempted only if it is known that a previous attempt definitely failed.

One exception to this rule are datagram invocations over UDP transports. For these, duplicated UDP packets can lead to a violation of at-most-once semantics.

At-most-once semantics are important because they guarantee that operations that are not *idempotent* can be used safely. An idempotent operation is an operation that, if executed twice, has the same effect as if executed once. For example, `x = 1;` is an idempotent operation: if we execute the operation twice, the end result is the same as if we had executed it once. On the other hand, `x++;` is not idempotent: if we execute the operation twice, the end result is not the same as if we had executed it once.

Without at-most-once semantics, we can build distributed systems that are more robust in the presence of network failures. However, realistic systems require non-idempotent operations, so at-most-once semantics are a necessity, even though they make the system less robust in the presence of network failures. Ice permits you to mark individual operations as idempotent. For such operations, the Ice run time

Copyright 2016, ZeroC, Inc.

uses a more aggressive error recovery mechanism than for non-idempotent operations.

Synchronous Method Invocation

By default, the request dispatch model used by Ice is a synchronous remote procedure call: an operation invocation behaves like a local procedure call, that is, the client thread is suspended for the duration of the call and resumes when the call completes (and all its results are available).

Asynchronous Method Invocation

Ice also supports *asynchronous method invocation (AMI)*: clients can invoke operations *asynchronously*, that is, the client uses a proxy as usual to invoke an operation but, in addition to passing the normal parameters, also passes a *callback object* and the client invocation returns immediately. Once the operation completes, the client-side run time invokes a method on the callback object passed initially, passing the results of the operation to the callback object (or, in case of failure, passing exception information).

The server cannot distinguish an asynchronous invocation from a synchronous one — either way, the server simply sees that a client has invoked an operation on an object.

Asynchronous Method Dispatch

Asynchronous method dispatch (AMD) is the server-side equivalent of AMI. For synchronous dispatch (the default), the server-side run time up-calls into the application code in the server in response to an operation invocation. While the operation is executing (or sleeping, for example, because it is waiting for data), a thread of execution is tied up in the server; that thread is released only when the operation completes.

With asynchronous method dispatch, the server-side application code is informed of the arrival of an operation invocation. However, instead of being forced to process the request immediately, the server-side application can choose to delay processing of the request and, in doing so, releases the execution thread for the request. The server-side application code is now free to do whatever it likes. Eventually, once the results of the operation are available, the server-side application code makes an API call to inform the server-side Ice run time that a request that was dispatched previously is now complete; at that point, the results of the operation are returned to the client.

Asynchronous method dispatch is useful if, for example, a server offers operations that block clients for an extended period of time. For example, the server may have an object with a `get` operation that returns data from an external, asynchronous data source and that blocks clients until the data becomes available. With synchronous dispatch, each client waiting for data to arrive ties up an execution thread in the server. Clearly, this approach does not scale beyond a few dozen clients. With asynchronous dispatch, hundreds or thousands of clients can be blocked in the same operation invocation without tying up any threads in the server.

Another way to use asynchronous method dispatch is to complete an operation, so the results of the operation are returned to the client, but to keep the execution thread of the operation beyond the duration of the operation invocation. This allows you to continue processing after results have been returned to the client, for example, to perform cleanup or write updates to persistent storage.

Synchronous and asynchronous method dispatch are transparent to the client, that is, the client cannot tell whether a server chose to process a request synchronously or asynchronously.

Oneway Method Invocation

Clients can invoke an operation as a *oneway* operation. A oneway invocation has "best effort" semantics. For a oneway invocation, the client-side run time hands the invocation to the local transport, and the invocation completes on the client side as soon as the local transport has buffered the invocation. The actual invocation is then sent asynchronously by the operating system. The server does not reply to oneway invocations, that is, traffic flows only from client to server, but not vice versa.

Oneway invocations are unreliable. For example, the target object may not exist, in which case the invocation is simply lost. Similarly, the operation may be dispatched to a servant in the server, but the operation may fail (for example, because parameter values are invalid); if so, the client receives no notification that something has gone wrong.

Oneway invocations are possible only on operations that do not have a return value, do not have out-parameters, and do not throw user exceptions.

To the application code on the server-side, oneway invocations are transparent, that is, there is no way to distinguish a twoway invocation from a oneway invocation.

Oneway invocations are available only if the target object offers a stream-oriented transport, such as TCP/IP or SSL.

Note that, even though oneway operations are sent over a stream-oriented transport, they may be processed out of order in the server. This can happen because each invocation may be dispatched in its own thread: even though the invocations are *initiated* in the order in which the invocations arrive at the server, this does not mean that they will be *processed* in that order — the vagaries of thread scheduling can result in a oneway invocation completing before other oneway invocations that were received earlier.

Copyright 2016, ZeroC, Inc.

Batched Oneway Method Invocation

Each oneway invocation sends a separate message to the server. For a series of short messages, the overhead of doing so is considerable: the client- and server-side run time each must switch between user mode and kernel mode for each message and, at the networking level, each message incurs the overheads of flow-control and acknowledgement.

Batched oneway invocations allow you to send a series of oneway invocations as a single message: every time you invoke a batched oneway operation, the invocation is buffered in the client-side run time. Once you have accumulated all the oneway invocations you want to send, you make a separate API call to send all the invocations at once. The client-side run time then sends all of the buffered invocations in a single message, and the server receives all of the invocations in a single message. This avoids the overhead of repeatedly trapping into the kernel for both client and server, and is much easier on the network between them because one large message can be transmitted more efficiently than many small ones.

The individual invocations in a batched oneway message are dispatched by a single thread in the order in which they were placed into the batch. This guarantees that the individual operations in a batched oneway message are processed in order in the server.

Batched oneway invocations are particularly useful for messaging services, such as IceStorm, and for fine-grained interfaces that offer `set` operations for small attributes.

Datagram Invocations

Datagram invocations have "best effort" semantics similar to oneway invocations. However, datagram invocations require the object to offer UDP as a transport (whereas oneway invocations require TCP/IP).

Like a oneway invocation, a datagram invocation can be made only if the operation does not have a return value, out-parameters, or user exceptions. A datagram invocation uses UDP to invoke the operation. The operation returns as soon as the local UDP stack has accepted the message; the actual operation invocation is sent asynchronously by the network stack behind the scenes.

Datagrams, like oneway invocations, are unreliable: the target object may not exist in the server, the server may not be running, or the operation may be invoked in the server but fail due to invalid parameters sent by the client. As for oneway invocations, the client receives no notification of such errors.

However, unlike oneway invocations, datagram invocations have a number of additional error scenarios:

- Individual invocations may simply be lost in the network.
  This is due to the unreliable delivery of UDP packets. For example, if you invoke three operations in sequence, the middle invocation may be lost. (The same thing cannot happen for oneway invocations — because they are delivered over a connection-oriented transport, individual invocations cannot be lost.)
- Individual invocations may arrive out of order.
  Again, this is due to the nature of UDP datagrams. Because each invocation is sent as a separate datagram, and individual datagrams can take different paths through the network, it can happen that invocations arrive in an order that differs from the order in which they were sent.

Datagram invocations are well suited for small messages on LANs, where the likelihood of loss is small. They are also suited to situations in which low latency is more important than reliability, such as for fast, interactive internet applications. Finally, datagram invocations can be used to multicast messages to multiple servers simultaneously.

Batched Datagram Invocations

As for batched oneway invocations, *batched datagram invocations* allow you to accumulate a number of invocations in a buffer and then send the entire buffer as a single datagram by making an API call to flush the buffer. Batched datagrams reduce the overhead of repeated system calls and allow the underlying network to operate more efficiently. However, batched datagram invocations are useful only for batched messages whose total size does not substantially exceed the PDU limit of the network: if the size of a batched datagram gets too large, UDP fragmentation makes it more likely that one or more fragments are lost, which results in the loss of the entire batched message. However, you are guaranteed that either all invocations in a batch will be delivered, or none will be delivered. It is impossible for individual invocations within a batch to be lost.

Batched datagrams use a single thread in the server to dispatch the individual invocations in a batch. This guarantees that the invocations are made in the order in which they were queued — invocations cannot appear to be reordered in the server.

Run-Time Exceptions

Any operation invocation can raise a *run-time exception*. Run-time exceptions are pre-defined by the Ice run time and cover common error conditions, such as connection failure, connection timeout, or resource allocation failure. Run-time exceptions are presented to the application as native exceptions and so integrate neatly with the native exception handling capabilities of languages that support exception handling.

Copyright 2016, ZeroC, Inc.

User Exceptions

A server indicates application-specific error conditions by raising *user exceptions* to clients. User exceptions can carry an arbitrary amount of complex data and can be arranged into inheritance hierarchies, which makes it easy for clients to handle categories of errors generically, by catching an exception that is further up the inheritance hierarchy. Like run-time exceptions, user exceptions map to native exceptions.

Properties

Much of the Ice run time is configurable via *properties*. Properties are name-value pairs, such as `Ice.Default.Protocol=tcp`. Properties are typically stored in text files and parsed by the Ice run time to configure various options, such as the thread pool size, the level of tracing, and various other configuration parameters.

See Also

- The Slice Language
- Proxies for Ice Objects
- Locators
- Object Life Cycle
- Bidirectional Connections
- Glacier2
- IceStorm

Copyright 2016, ZeroC, Inc.

Slice (Specification Language for Ice)

Each Ice object has an interface with a number of operations. Interfaces, operations, and the types of data that are exchanged between client and server are defined using the *Slice language*. Slice allows you to define the client-server contract in a way that is independent of a specific programming language, such as C++, Java, or C#. The Slice definitions are compiled by a compiler into an API for a specific programming language, that is, the part of the API that is specific to the interfaces and types you have defined consists of generated code.

See Also

- The Slice Language

Copyright 2016, ZeroC, Inc.

Overview of the Language Mappings

The rules that govern how each Slice construct is translated into a specific programming language are known as *language mappings*. For example, for the C++ mapping, a Slice sequence appears as a `std::vector`, whereas, for the Java mapping, a Slice sequence appears as a Java array. In order to determine what the API for a specific Slice construct looks like, you only need the Slice definition and knowledge of the language mapping rules. The rules are simple and regular enough to make it unnecessary to read the generated code to work out how to use the generated API.

Of course, you are free to peruse the generated code. However, as a rule, that is inefficient because the generated code is not necessarily suitable for human consumption. We recommend that you familiarize yourself with the language mapping rules; that way, you can mostly ignore the generated code and need to refer to it only when you are interested in some specific detail.

Currently, Ice provides language mappings for C++, C#, Java, JavaScript, Python, Objective-C, and, for the client side, PHP and Ruby.

See Also

- C++ Mapping
- C# Mapping
- Java Mapping
- JavaScript Mapping
- Objective-C Mapping
- PHP Mapping
- Python Mapping
- Ruby Mapping

Copyright 2016, ZeroC, Inc.

Client and Server Structure

Ice clients and servers have the logical internal structure:

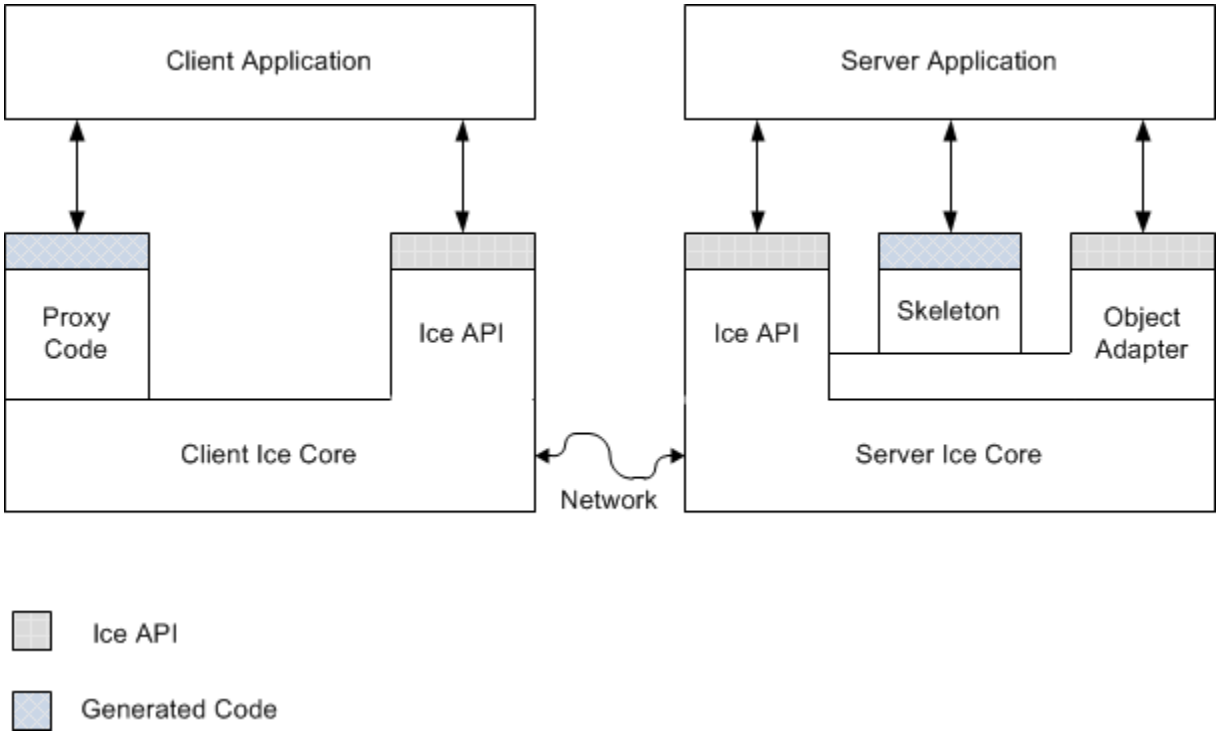

Ice Client and Server Structure

Both client and server consist of a mixture of application code, library code, and code generated from Slice definitions:

- The Ice core contains the client- and server-side run-time support for remote communication. Much of this code is concerned with the details of networking, threading, byte ordering, and many other networking-related issues that we want to keep away from application code. The Ice core is provided as a number of libraries that client and server use.
- The generic part of the Ice core (that is, the part that is independent of the specific types you have defined in Slice) is accessed through the Ice API. You use the Ice API to take care of administrative chores, such as initializing and finalizing the Ice run time. The Ice API is identical for clients and servers (although servers use a larger part of the API than clients).
- The proxy code is generated from your Slice definitions and, therefore, specific to the types of objects and data you have defined in Slice. The proxy code has two major functions:
  - It provides a down-call interface for the client. Calling a function in the generated proxy API ultimately ends up sending an RPC message to the server that invokes a corresponding function on the target object.
  - It provides *marshaling* and *unmarshaling* code. Marshaling is the process of serializing a complex data structure, such as a sequence or a dictionary, for transmission on the wire. The marshaling code converts data into a form that is standardized for transmission and independent of the endian-ness and padding rules of the local machine. Unmarshaling is the reverse of marshaling, that is, deserializing data that arrives over the network and reconstructing a local representation of the data in types that are appropriate for the programming language in use.
- The skeleton code is also generated from your Slice definition and, therefore, specific to the types of objects and data you have defined in Slice. The skeleton code is the server-side equivalent of the client-side proxy code: it provides an up-call interface that permits the Ice run time to transfer the thread of control to the application code you write. The skeleton also contains marshaling and unmarshaling code, so the server can receive parameters sent by the client, and return parameters and exceptions to the client.
- The object adapter is a part of the Ice API that is specific to the server side: only servers use object adapters. An object adapter has several functions:
  - The object adapter maps incoming requests from clients to specific methods on programming-language objects. In other words, the object adapter tracks which servants with what object identity are in memory.
  - The object adapter is associated with one or more transport endpoints. If more than one transport endpoint is associated with an adapter, the servants incarnating objects within the adapter can be reached via multiple transports. For example, you can associate both a TCP/IP and a UDP endpoint with an adapter, to provide alternate quality-of-service and performance characteristics.
  - The object adapter is responsible for the creation of proxies that can be passed to clients. The object adapter knows about

Copyright 2016, ZeroC, Inc.

the type, identity, and transport details of each of its objects and embeds the correct details when the server-side application code requests the creation of a proxy.

Note that, as far as the process view is concerned, there are only two processes involved: the client and the server. All the run time support for distributed communication is provided by the Ice libraries and the code that is generated from Slice definitions. (For indirect proxies, a location service is required to resolve proxies to transport endpoints.)

See Also

- Hello World Application
- Locators

Copyright 2016, ZeroC, Inc.

Overview of the Ice Protocol

Ice provides an RPC protocol that can use either TCP/IP or UDP as an underlying transport. In addition, Ice also allows you to use SSL as a transport, so all communication between client and server is encrypted.

The Ice protocol defines:

- a number of message types, such as request and reply message types,
- a protocol state machine that determines in what sequence different message types are exchanged by client and server, together with the associated connection establishment and tear-down semantics for TCP/IP,
- encoding rules that determine how each type of data is represented on the wire,
- a header for each message type that contains details such as the message type, the message size, and the protocol and encoding version in use.

Ice also supports compression on the wire: by setting a configuration parameter, you can arrange for all network traffic to be compressed to conserve bandwidth. This is useful if your application exchanges large amounts of data between client and server.

The Ice protocol is suitable for building highly-efficient event forwarding mechanisms because it permits forwarding of a message without knowledge of the details of the information inside a message. This means that messaging switches need not do any unmarshaling and remarshaling of messages — they can forward a message by simply treating it as an opaque buffer of bytes.

The Ice protocol also supports bidirectional operation: if a server wants to send a message to a callback object provided by the client, the callback can be made over the connection that was originally created by the client. This feature is especially important when the client is behind a firewall that permits outgoing connections, but not incoming connections.

See Also

- The Ice Protocol
- IceSSL
- Bidirectional Connections

Copyright 2016, ZeroC, Inc.

Ice Services Overview

The Ice core provides a sophisticated client-server platform for distributed application development. However, realistic applications usually require more than just a remoting capability: typically, you also need a way to start servers on demand, distribute proxies to clients, distribute asynchronous events, configure your application, distribute patches for an application, and so on.

Ice ships with a number of services that provide these and other features. The services are implemented as Ice servers to which your application acts as a client. None of the services use Ice-internal features that are hidden from application developers so, in theory, you could develop equivalent services yourself. However, having these services available as part of the platform allows you to focus on application development instead of having to build a lot of infrastructure first. Moreover, building such services is not a trivial effort, so it pays to know what is available and use it instead of reinventing your own wheel.

On this page:

- Freeze and FreezeScript
- IceGrid Service
- IceBox Server
- IceStorm
- IcePatch2
- Glacier2

Freeze and FreezeScript

Ice has a built-in object persistence service, known as *Freeze*. Freeze makes it easy to store object state in a database: you define the state stored by your objects in Slice, and the Freeze compiler generates code that stores and retrieves object state to and from a database. Freeze uses Berkeley DB as its database.

Ice also offers a tool set collectively called FreezeScript that makes it easier to maintain databases and to migrate the contents of existing databases to a new schema if the type definitions of objects change.

IceGrid Service

IceGrid is an implementation of an Ice location service that resolves the symbolic information in an indirect proxy to a protocol-address pair for indirect binding. A location service is only the beginning of IceGrid's capabilities.

IceGrid:

- allows you to register servers for automatic start-up: instead of requiring a server to be running at the time a client issues a request, IceGrid starts servers on demand, when the first client request arrives.
- provides tools that make it easy to configure complex applications containing several servers.
- supports replication and load-balancing.
- automates the distribution and patching of server executables and dependent files.
- provides a simple query service that allows clients to obtain proxies for objects they are interested in.

IceBox Server

IceBox is a simple application server that can orchestrate the starting and stopping of a number of application components. Application components can be deployed as a dynamic library instead of as a process. This reduces overall system load, for example, by allowing you to run several application components in a single Java virtual machine instead of having multiple processes, each with its own virtual machine.

IceStorm

IceStorm is a publish-subscribe service that decouples clients and servers. Fundamentally, IceStorm acts as a distribution switch for events. Publishers send events to the service, which, in turn, passes the events to subscribers. In this way, a single event published by a publisher can be sent to multiple subscribers. Events are categorized by topic, and subscribers specify the topics they are interested in. Only events that match a subscriber's topic are sent to that subscriber. The service permits selection of a number of quality-of-service criteria to allow applications to choose the appropriate trade-off between reliability and performance.

IceStorm is particularly useful if you have a need to distribute information to large numbers of application components. (A typical example is a stock ticker application with a large number of subscribers.) IceStorm decouples the publishers of information from subscribers and takes care of the redistribution of the published events. In addition, IceStorm can be run as a *federated* service, that is, multiple instances of the service can be run on different machines to spread the processing load over a number of CPUs.

IcePatch2

Copyright 2016, ZeroC, Inc.

IcePatch2 is a software patching service. It allows you to easily distribute software updates to clients. Clients simply connect to the IcePatch2 server and request updates for a particular application. The service automatically checks the version of the client's software and downloads any updated application components in a compressed format to conserve bandwidth. Software patches can be secured using the Glacier2 service, so only authorized clients can download software updates.

IcePatch2 supersedes IcePatch, which was a previous version of this service.

Glacier2

Glacier2 is the Ice firewall traversal service: it allows clients and servers to securely communicate through a firewall without compromising security. Client-server traffic is SSL-encrypted using public key certificates and is bidirectional. Glacier2 offers support for mutual authentication as well as secure session management.

Glacier2 supersedes Glacier, which was a previous version of this service

See Also

- IceGrid
- Freeze
- FreezeScript
- Glacier2
- IceBox
- IceStorm
- IcePatch2

Copyright 2016, ZeroC, Inc.

Architectural Benefits of Ice

The Ice architecture provides a number of benefits to application developers:

- Object-oriented semantics
  Ice fully preserves the object-oriented paradigm "across the wire." All operation invocations use late binding, so the implementation of an operation is chosen depending on the actual run-time (not static) type of an object.
- Support for synchronous and asynchronous messaging
  Ice provides both synchronous and asynchronous operation invocation and dispatch, as well as publish-subscribe messaging via IceStorm. This allows you to choose a communication model according to the needs of your application instead of having to shoe-horn the application to fit a single model.
- Support for multiple interfaces
  With facets, objects can provide multiple, unrelated interfaces while retaining a single object identity across these interfaces. This provides great flexibility, particularly as an application evolves but needs to remain compatible with older, already deployed clients.
- Machine independence
  Clients and servers are shielded form idiosyncrasies of the underlying machine architecture. Issues such as byte ordering and padding are hidden from application code.
- Language independence
  Client and server can be developed independently and in different programming languages. The Slice definition used by both client and server establishes the interface contract between them and is the only thing they need to agree on.
- Implementation independence
  Clients are unaware of how servers implement their objects. This means that the implementation of a server can be changed after clients are deployed, for example, to use a different persistence mechanism or even a different programming language.
- Operating system independence
  The Ice APIs are fully portable, so the same source code compiles and runs under both Windows and Unix.
- Threading support
  The Ice run time is fully threaded and APIs are thread-safe. No effort (beyond synchronizing access to shared data) is required on part of the application developer to develop threaded, high-performance clients and servers.
- Transport independence
  Ice currently offers both TCP/IP and UDP as transport protocols. Neither client nor server code are aware of the underlying transport. (The desired transport can be chosen by a configuration parameter.)
- Location and server transparency
  The Ice run time takes care of locating objects and managing the underlying transport mechanism, such as opening and closing connections. Interactions between client and server appear connection-less. Via IceGrid, you can arrange for servers to be started on demand if they are not running at the time a client invokes an operation. Servers can be migrated to different physical addresses without breaking proxies held by clients, and clients are completely unaware how object implementations are distributed over server processes.
- Security
  Communications between client and server can be fully secured with strong encryption over SSL, so applications can use unsecured public networks to communicate securely. Via Glacier2, you can implement secure forwarding of requests through a firewall, with full support for callbacks.
- Built-in persistence
  With Freeze, creating persistent object implementations becomes trivial. Ice comes with built-in support for Berkeley DB, which is a high-performance database.
- Source code availability
  The source code for Ice is available. While it is not necessary to have access to the source code to use the platform, it allows you to see how things are implemented or port the code to a new operating system.

Overall, Ice provides a state-of-the art development and deployment environment for distributed computing that is more complete than any other platform we are aware of.

See Also

- Ice Architecture
- Ice Services Overview

Copyright 2016, ZeroC, Inc.

Hello World Application

This section presents a very simple (but complete) client and server.

Writing an Ice application involves the following steps:

1. Write a Slice definition and compile it.
2. Write a server and compile it.
3. Write a client and compile it.

If someone else has written the server already and you are only writing a client, you do not need to write the Slice definition, only compile it (and, obviously, you do not need to write the server in that case).

The application described here enables remote printing: a client sends the text to be printed to a server, which in turn sends that text to a printer. For simplicity (and because we do not want to concern ourselves with the idiosyncrasies of print spoolers for various platforms), our printer will simply print to a terminal instead of a real printer. This is no great loss: the purpose of the exercise is to show how a client can communicate with a server; once the thread of control has reached the server application code, that code can of course do anything it likes (including sending the text to a real printer). How to do this is independent of Ice and therefore not relevant here.

Much of the detail of the source code will remain unexplained for now. The intent is to show you how to get started and give you a feel for what the development environment looks like; we will provide all the detail throughout the remainder of this manual.

Topics

- Writing a Slice Definition
- Writing an Ice Application with C++
- Writing an Ice Application with C-Sharp
- Writing an Ice Application with Java
- Writing an Ice Application with JavaScript
- Writing an Ice Application with Objective-C
- Writing an Ice Application with PHP
- Writing an Ice Application with Python
- Writing an Ice Application with Ruby
- Writing an Ice Application with Visual Basic

Copyright 2016, ZeroC, Inc.

Writing a Slice Definition

The first step in writing any Ice application is to write a Slice definition containing the interfaces that are used by the application. For our minimal printing application, we write the following Slice definition:

Slice

```
module Demo {
    interface Printer {
        void printString(string s);
    };
};
```

We save this text in a file called `Printer.ice`.

Our Slice definitions consist of the module `Demo` containing a single interface called `Printer`. For now, the interface is very simple and provides only a single operation, called `printString`. The `printString` operation accepts a string as its sole input parameter; the text of that string is what appears on the (possibly remote) printer.

See Also

- The Slice Language

Copyright 2016, ZeroC, Inc.

Writing an Ice Application with C++

This page shows how to create an Ice application with C++.

On this page:

- Compiling a Slice Definition for C++
- Writing and Compiling a Server in C++
- Writing and Compiling a Client in C++
- Running Client and Server in C++

Compiling a Slice Definition for C++

The first step in creating our C++ application is to compile our Slice definition to generate C++ proxies and skeletons. You can compile the definition as follows:

```
$ slice2cpp Printer.ice
```

The `slice2cpp` compiler produces two C++ source files from this definition, `Printer.h` and `Printer.cpp`.

- `Printer.h`
  The `Printer.h` header file contains C++ type definitions that correspond to the Slice definitions for our `Printer` interface. This header file must be included in both the client and the server source code.
- `Printer.cpp`
  The `Printer.cpp` file contains the source code for our `Printer` interface. The generated source contains type-specific run-time support for both clients and servers. For example, it contains code that marshals parameter data (the string passed to the `printString` operation) on the client side and unmarshals that data on the server side.
  The `Printer.cpp` file must be compiled and linked into both client and server.

Writing and Compiling a Server in C++

The source code for the server takes only a few lines and is shown in full here:

Copyright 2016, ZeroC, Inc.

C++

```
#include <Ice/Ice.h>
#include <Printer.h>

using namespace std;
using namespace Demo;

class PrinterI : public Printer {
public:
    virtual void printString(const string& s, const Ice::Current&);
};

void
PrinterI::
printString(const string& s, const Ice::Current&)
{
    cout << s << endl;
}

int
main(int argc, char* argv[])
{
    int status = 0;
    Ice::CommunicatorPtr ic;
    try {
        ic = Ice::initialize(argc, argv);
        Ice::ObjectAdapterPtr adapter =

ic->createObjectAdapterWithEndpoints("SimplePrinterAdapter", "default -p 10000");
        Ice::ObjectPtr object = new PrinterI;
        adapter->add(object, ic->stringToIdentity("SimplePrinter"));
        adapter->activate();
        ic->waitForShutdown();
    } catch (const Ice::Exception& e) {
        cerr << e << endl;
        status = 1;
    } catch (const char* msg) {
        cerr << msg << endl;
        status = 1;
    }
    if (ic) {
        try {
            ic->destroy();
        } catch (const Ice::Exception& e) {
            cerr << e << endl;
            status = 1;
        }
    }
    return status;
}
```

There appears to be a lot of code here for something as simple as a server that just prints a string. Do not be concerned by this: most of the

Copyright 2016, ZeroC, Inc.

preceding code is boiler plate that never changes. For this very simple server, the code is dominated by this boiler plate.

Every Ice source file starts with an include directive for `Ice.h`, which contains the definitions for the Ice run time. We also include `Printer.h`, which was generated by the Slice compiler and contains the C++ definitions for our printer interface, and we import the contents of the `std` and `Demo` namespaces for brevity in the code that follows:

C++

```
#include <Ice/Ice.h>
#include <Printer.h>

using namespace std;
using namespace Demo;
```

Our server implements a single printer servant, of type `PrinterI`. Looking at the generated code in `Printer.h`, we find the following (tidied up a little to get rid of irrelevant detail):

C++

```
namespace Demo {
    class Printer : virtual public Ice::Object {
    public:
        virtual void printString(const std::string&, const Ice::Current&
= Ice::Current()) = 0;
    };
};
```

The `Printer` skeleton class definition is generated by the Slice compiler. (Note that the `printString` method is pure virtual so the skeleton class cannot be instantiated.) Our servant class inherits from the skeleton class to provide an implementation of the pure virtual `printString` method. (By convention, we use an `I`-suffix to indicate that the class implements an interface.)

C++

```
class PrinterI : public Printer {
public:
    virtual void printString(const string& s, const Ice::Current&);
};
```

The implementation of the `printString` method is trivial: it simply writes its string argument to `stdout`:

C++

```
void
PrinterI::
printString(const string& s, const Ice::Current&)
{
    cout << s << endl;
}
```

Note that `printString` has a second parameter of type `Ice::Current`. As you can see from the definition of `Printer::printString`,

Copyright 2016, ZeroC, Inc.

the Slice compiler generates a default argument for this parameter, so we can leave it unused in our implementation. (We will examine the purpose of the `Ice::Current` parameter later.)

What follows is the server main program. Note the general structure of the code:

C++

```
int
main(int argc, char* argv[])
{
    int status = 0;
    Ice::CommunicatorPtr ic;
    try {

        // Server implementation here...

    } catch (const Ice::Exception& e) {
        cerr << e << endl;
        status = 1;
    } catch (const char* msg) {
        cerr << msg << endl;
        status = 1;
    }
    if (ic) {
        try {
            ic->destroy();
        } catch (const Ice::Exception& e) {
            cerr << e << endl;
            status = 1;
        }
    }
    return status;
}
```

The body of `main` contains the declaration of two variables, `status` and `ic`. The `status` variable contains the exit status of the program and the `ic` variable, of type `Ice::CommunicatorPtr`, contains the main handle to the Ice run time.

Following these declarations is a `try` block in which we place all the server code, followed by two `catch` handlers. The first handler catches all exceptions that may be thrown by the Ice run time; the intent is that, if the code encounters an unexpected Ice run-time exception anywhere, the stack is unwound all the way back to `main`, which prints the exception and then returns failure to the operating system. The second handler catches string constants; the intent is that, if we encounter a fatal error condition somewhere in our code, we can simply throw a string literal with an error message. Again, this unwinds the stack all the way back to `main`, which prints the error message and then returns failure to the operating system.

Following the `try` block, we see a bit of cleanup code that calls the `destroy` method on the communicator (provided that the communicator was initialized). The cleanup call is outside the first `try` block for a reason: we must ensure that the Ice run time is finalized whether the code terminates normally or terminates due to an exception.

Failure to call `destroy` on the communicator before the program exits results in undefined behavior.

The body of the first `try` block contains the actual server code:

Copyright 2016, ZeroC, Inc.

C++

```
        ic = Ice::initialize(argc, argv);
        Ice::ObjectAdapterPtr adapter =

ic->createObjectAdapterWithEndpoints("SimplePrinterAdapter", "default -p 10000");
        Ice::ObjectPtr object = new PrinterI;
        adapter->add(object, ic->stringToIdentity("SimplePrinter"));
        adapter->activate();
        ic->waitForShutdown();
```

The code goes through the following steps:

1. We initialize the Ice run time by calling `Ice::initialize`. (We pass `argc` and `argv` to this call because the server may have command-line arguments that are of interest to the run time; for this example, the server does not require any command-line arguments.) The call to `initialize` returns a smart pointer to an `Ice::Communicator` object, which is the main object in the Ice run time.
2. We create an object adapter by calling `createObjectAdapterWithEndpoints` on the `Communicator` instance. The arguments we pass are `"SimplePrinterAdapter"` (which is the name of the adapter) and `"default -p 10000"`, which instructs the adapter to listen for incoming requests using the default protocol (TCP/IP) at port number 10000.
3. At this point, the server-side run time is initialized and we create a servant for our `Printer` interface by instantiating a `PrinterI` object.
4. We inform the object adapter of the presence of a new servant by calling `add` on the adapter; the arguments to `add` are the servant we have just instantiated, plus an identifier. In this case, the string `"SimplePrinter"` is the name of the Ice object. (If we had multiple printers, each would have a different name or, more correctly, a different *object identity*.)
5. Next, we activate the adapter by calling its `activate` method. (The adapter is initially created in a holding state; this is useful if we have many servants that share the same adapter and do not want requests to be processed until after all the servants have been instantiated.) The server starts to process incoming requests from clients as soon as the adapter is activated.
6. Finally, we call `waitForShutdown`. This call suspends the calling thread until the server implementation terminates, either by making a call to shut down the run time, or in response to a signal. (For now, we will simply interrupt the server on the command line when we no longer need it.)

Note that, even though there is quite a bit of code here, that code is essentially the same for all servers. You can put that code into a helper class and, thereafter, will not have to bother with it again. (Ice provides such a helper class, called `Ice::Application`.) As far as actual application code is concerned, the server contains only a few lines: six lines for the definition of the `PrinterI` class, plus three lines to instantiate a `PrinterI` object and register it with the object adapter.

Assuming that we have the server code in a file called `Server.cpp`, we can compile it as follows:

```
$ c++ -I. -c Printer.cpp Server.cpp
```

This compiles both our application code and the code that was generated by the Slice compiler. Depending on your platform, you may have to add additional include directives or other options to the compiler; please see the demo programs that ship with Ice for the details.

Finally, we need to link the server into an executable:

```
$ c++ -o server Printer.o Server.o -lIce -lIceUtil
```

Again, depending on the platform, the actual list of libraries you need to link against may be longer. The demo programs that ship with Ice contain all the detail. The important point to note here is that the Ice run time is shipped in two libraries, `libIce` and `libIceUtil`.

Writing and Compiling a Client in C++

The client code looks very similar to the server. Here it is in full:

Copyright 2016, ZeroC, Inc.

C++

```
#include <Ice/Ice.h>
#include <Printer.h>

using namespace std;
using namespace Demo;

int
main(int argc, char* argv[])
{
    int status = 0;
    Ice::CommunicatorPtr ic;
    try {
        ic = Ice::initialize(argc, argv);
        Ice::ObjectPrx base = ic->stringToProxy("SimplePrinter:default -p 10000");
        PrinterPrx printer = PrinterPrx::checkedCast(base);
        if (!printer)
            throw "Invalid proxy";

        printer->printString("Hello World!");
    } catch (const Ice::Exception& ex) {
        cerr << ex << endl;
        status = 1;
    } catch (const char* msg) {
        cerr << msg << endl;
        status = 1;
    }
    if (ic)
        ic->destroy();
    return status;
}
```

Note that the overall code layout is the same as for the server: we include the headers for the Ice run time and the header generated by the Slice compiler, and we use the same `try` block and `catch` handlers to deal with errors.

The code in the `try` block does the following:

1. As for the server, we initialize the Ice run time by calling `Ice::initialize`.
2. The next step is to obtain a proxy for the remote printer. We create a proxy by calling `stringToProxy` on the communicator, with the string `"SimplePrinter:default -p 10000"`. Note that the string contains the object identity and the port number that were used by the server. (Obviously, hard-coding object identities and port numbers into our applications is a bad idea, but it will do for now; we will see more architecturally sound ways of doing this when we discuss IceGrid.)
3. The proxy returned by `stringToProxy` is of type `Ice::ObjectPrx`, which is at the root of the inheritance tree for interfaces and classes. But to actually talk to our printer, we need a proxy for a `Printer` interface, not an `Object` interface. To do this, we need to do a down-cast by calling `PrinterPrx::checkedCast`. A checked cast sends a message to the server, effectively asking "is this a proxy for a `Printer` interface?" If so, the call returns a proxy to a `Printer`; otherwise, if the proxy denotes an interface of some other type, the call returns a null proxy.
4. We test that the down-cast succeeded and, if not, throw an error message that terminates the client.
5. We now have a live proxy in our address space and can call the `printString` method, passing it the time-honored `"Hello World!"` string. The server prints that string on its terminal.

Compiling and linking the client looks much the same as for the server:

Copyright 2016, ZeroC, Inc.

```
$ c++ -I. -I$ICE_HOME/include -c Printer.cpp Client.cpp
$ c++ -o client Printer.o Client.o -lIce -lIceUtil
```

Running Client and Server in C++

To run client and server, we first start the server in a separate window:

```
$ ./server
```

At this point, we won't see anything because the server simply waits for a client to connect to it. We run the client in a different window:

```
$ ./client
$
```

The client runs and exits without producing any output; however, in the server window, we see the `"Hello World!"` that is produced by the printer. To get rid of the server, we interrupt it on the command line for now. (We will see cleaner ways to terminate a server in our discussion of `Ice::Application`.)

If anything goes wrong, the client will print an error message. For example, if we run the client without having first started the server, we get:

```
Network.cpp:471: Ice::ConnectFailedException:
connect failed: Connection refused
```

See Also

- Client-Side Slice-to-C++ Mapping
- Server-Side Slice-to-C++ Mapping
- The `Ice::Application` Class
- The Current Object
- IceGrid

Copyright 2016, ZeroC, Inc.

Writing an Ice Application with C-Sharp

This page shows how to create an Ice application with C#.

On this page:

- Compiling a Slice Definition for C#
- Writing and Compiling a Server in C#
- Writing and Compiling a Client in C#
- Running Client and Server in C#

Compiling a Slice Definition for C#

The first step in creating our C# application is to compile our Slice definition to generate C# proxies and skeletons. You can compile the definition as follows:

```
> mkdir generated
> slice2cs --output-dir generated Printer.ice
```

The `--output-dir` option instructs the compiler to place the generated files into the `generated` directory. This avoids cluttering the working directory with the generated files. The `slice2cs` compiler produces a single source file, `Printer.cs`, from this definition. The exact contents of this file do not concern us for now — it contains the generated code that corresponds to the `Printer` interface we defined in `Printer.ice`.

Writing and Compiling a Server in C#

To implement our `Printer` interface, we must create a servant class. By convention, a servant class uses the name of its interface with an `I`-suffix, so our servant class is called `PrinterI` and placed into a source file `Server.cs`:

C#

```
using System;

public class PrinterI : Demo.PrinterDisp_
{
    public override void printString(string s, Ice.Current current)
    {
        Console.WriteLine(s);
    }
}
```

The `PrinterI` class inherits from a base class called `PrinterDisp_`, which is generated by the `slice2cs` compiler. The base class is abstract and contains a `printString` method that accepts a string for the printer to print and a parameter of type `Ice.Current`. (For now we will ignore the `Ice.Current` parameter.) Our implementation of the `printString` method simply writes its argument to the terminal.

The remainder of the server code follows in `Server.cs` and is shown in full here:

Copyright 2016, ZeroC, Inc.

C#

```
public class Server
{
    public static void Main(string[] args)
    {
        int status = 0;
        Ice.Communicator ic = null;
        try {
            ic = Ice.Util.initialize(ref args);
            Ice.ObjectAdapter adapter =
                ic.createObjectAdapterWithEndpoints("SimplePrinterAdapter", "defaul
t -p 10000");
            Ice.Object obj = new PrinterI();
            adapter.add(obj, ic.stringToIdentity("SimplePrinter"));
            adapter.activate();
            ic.waitForShutdown();
        } catch (Exception e) {
            Console.Error.WriteLine(e);
            status = 1;
        }
        if (ic != null) {
            // Clean up
            //
            try {
                ic.destroy();
            } catch (Exception e) {
                Console.Error.WriteLine(e);
                status = 1;
            }
        }
        Environment.Exit(status);
    }
}
```

Note the general structure of the code:

Copyright 2016, ZeroC, Inc.

C#

```
public class Server
{
    public static void Main(string[] args)
    {
        int status = 0;
        Ice.Communicator ic = null;
        try {

            // Server implementation here...

        } catch (Exception e) {
            Console.Error.WriteLine(e);
            status = 1;
        }
        if (ic != null) {
            // Clean up
            //
            try {
                ic.destroy();
            } catch (Exception e) {
                Console.Error.WriteLine(e);
                status = 1;
            }
        }
        Environment.Exit(status);
    }
}
```

The body of `Main` contains a `try` block in which we place all the server code, followed by a `catch` block. The catch block catches all exceptions that may be thrown by the code; the intent is that, if the code encounters an unexpected run-time exception anywhere, the stack is unwound all the way back to `Main`, which prints the exception and then returns failure to the operating system.

Before the code exits, it destroys the communicator (if one was created successfully). Doing this is essential in order to correctly finalize the Ice run time: the program *must* call `destroy` on any communicator it has created; otherwise, undefined behavior results.

The structure shown above closely resembles that of other Ice language mappings such as C++ and Java. We can simplify things a bit in C# by taking advantage of the fact that the `Ice.Communicator` object implements `IDisposable`, which allows us to write a `using` statement as follows:

Copyright 2016, ZeroC, Inc.

C#

```
public class Server
{
    public static void Main(string[] args)
    {
        int status = 0;
        try {
            using(Ice.Communicator ic = Ice.Util.initialize(ref args)) {
                Ice.ObjectAdapter adapter =

    ic.createObjectAdapterWithEndpoints("SimplePrinterAdapter", "default -p 10000")
;
                Ice.Object obj = new PrinterI();
                adapter.add(obj, ic.stringToIdentity("SimplePrinter"));
                adapter.activate();
                ic.waitForShutdown();
            }
        } catch (Exception e) {
            Console.Error.WriteLine(e);
            status = 1;
        }
        Environment.Exit(status);
    }
}
```

We've eliminated the need for the clean-up code block because the `using` statement automatically calls `Dispose` on the communicator, which is equivalent to invoking its `destroy` method.

The body of our `try` block contains the actual server code:

C#

```
using(Ice.Communicator ic = Ice.Util.initialize(ref args)) {
                Ice.ObjectAdapter adapter =
                    ic.createObjectAdapterWithEndpoints("SimplePrinterAdapter",
"default -p 10000");
                Ice.Object obj = new PrinterI();
                adapter.add(obj, ic.stringToIdentity("SimplePrinter"));
                adapter.activate();
                ic.waitForShutdown();
            }
```

The code goes through the following steps:

1. We initialize the Ice run time by calling `Ice.Util.initialize`. (We pass `args` to this call because the server may have command-line arguments that are of interest to the run time; for this example, the server does not require any command-line arguments.) The call to `initialize` returns an `Ice.Communicator` reference, which is the main object in the Ice run time.
2. We create an object adapter by calling `createObjectAdapterWithEndpoints` on the `Communicator` instance. The arguments we pass are `"SimplePrinterAdapter"` (which is the name of the adapter) and `"default -p 10000"`, which instructs the adapter to listen for incoming requests using the default protocol (TCP/IP) at port number 10000.
3. At this point, the server-side run time is initialized and we create a servant for our `Printer` interface by instantiating a `PrinterI` object.

Copyright 2016, ZeroC, Inc.

4. We inform the object adapter of the presence of a new servant by calling `add` on the adapter; the arguments to `add` are the servant we have just instantiated, plus an identifier. In this case, the string `"SimplePrinter"` is the name of the Ice object. (If we had multiple printers, each would have a different name or, more correctly, a different *object identity*.)
5. Next, we activate the adapter by calling its `activate` method. (The adapter is initially created in a holding state; this is useful if we have many servants that share the same adapter and do not want requests to be processed until after all the servants have been instantiated.)
6. Finally, we call `waitForShutdown`. This call suspends the calling thread until the server implementation terminates, either by making a call to shut down the run time, or in response to a signal. (For now, we will simply interrupt the server on the command line when we no longer need it.)

Note that, even though there is quite a bit of code here, that code is essentially the same for all servers. You can put that code into a helper class and, thereafter, will not have to bother with it again. (Ice provides such a helper class, called `Ice.Application`.) As far as actual application code is concerned, the server contains only a few lines: seven lines for the definition of the `PrinterI` class, plus three lines to instantiate a `PrinterI` object and register it with the object adapter.

We can compile the server code as follows:

```
$ csc /reference:Ice.dll /lib:%ICE_HOME%\Assemblies Server.cs generated\Printer.cs
```

This compiles both our application code and the code that was generated by the Slice compiler. We assume that the `ICE_HOME` environment variable is set to the top-level directory containing the Ice run time. (For example, if you have installed Ice in `C:\Ice`, set `ICE_HOME` to that path.)

Writing and Compiling a Client in C#

The client code, in `Client.cs`, looks very similar to the server.

Here it is in full:

C#

```
using System;
using Demo;

public class Client
{
    public static void Main(string[] args)
    {
        int status = 0;
        try {
            using(Ice.Communicator ic = Ice.Util.initialize(ref args)) {

Ice.ObjectPrx obj = ic.stringToProxy("SimplePrinter:default -p 10000");
                PrinterPrx printer = PrinterPrxHelper.checkedCast(obj);
                if (printer == null)
                    throw new ApplicationException("Invalid proxy");

                printer.printString("Hello World!");
            }
        } catch (Exception e) {
            Console.Error.WriteLine(e);
            status = 1;
        }
        Environment.Exit(status);
    }
}
```


Copyright 2016, ZeroC, Inc.

Note that the overall code layout is the same as for the server: we use the same `try` and `catch` blocks to deal with errors. The code in the `try` block does the following:

1. As for the server, we initialize the Ice run time by calling `Ice.Util.initialize`.
2. The next step is to obtain a proxy for the remote printer. We create a proxy by calling `stringToProxy` on the communicator, with the string `"SimplePrinter:default -p 10000"`. Note that the string contains the object identity and the port number that were used by the server. (Obviously, hard-coding object identities and port numbers into our applications is a bad idea, but it will do for now; we will see more architecturally sound ways of doing this when we discuss IceGrid.
3. The proxy returned by `stringToProxy` is of type `Ice.ObjectPrx`, which is at the root of the inheritance tree for interfaces and classes. But to actually talk to our printer, we need a proxy for a `Printer` interface, not an `Object` interface. To do this, we need to do a down-cast by calling `PrinterPrxHelper.checkedCast`. A checked cast sends a message to the server, effectively asking "is this a proxy for a `Printer` interface?" If so, the call returns a proxy of type `Demo::Printer`; otherwise, if the proxy denotes an interface of some other type, the call returns null.
4. We test that the down-cast succeeded and, if not, throw an error message that terminates the client.
5. We now have a live proxy in our address space and can call the `printString` method, passing it the time-honored `"Hello World!"` string. The server prints that string on its terminal.

Compiling the client looks much the same as for the server:

```
$ csc /reference:Ice.dll /lib:%ICE_HOME%\Assemblies Client.cs generated\Printer.cs
```

Running Client and Server in C#

To run client and server, we first start the server in a separate window:

```
$ server.exe
```

At this point, we won't see anything because the server simply waits for a client to connect to it. We run the client in a different window:

```
$ client.exe
$
```

The client runs and exits without producing any output; however, in the server window, we see the `"Hello World!"` that is produced by the printer. To get rid of the server, we interrupt it on the command line for now. (We will see cleaner ways to terminate a server in our discussion of `Ice.Application`.)

If anything goes wrong, the client will print an error message. For example, if we run the client without having first started the server, we get something like the following:

```
Ice.ConnectionRefusedException
    error = 0
   at IceInternal.ProxyFactory.checkRetryAfterException(LocalException ex, Referenc
e ref, Int32 cnt)
   at Ice.ObjectPrxHelperBase.handleException__(ObjectDel_ delegate, LocalException
 ex, Int32 cnt)
   at Ice.ObjectPrxHelperBase.ice_isA(String id__, Dictionary`2 context__, Boolean
explicitContext__)
   at Ice.ObjectPrxHelperBase.ice_isA(String id__)
   at Demo.PrinterPrxHelper.checkedCast(ObjectPrx b)
   at Client.Main(String[] args)
Caused by: System.ComponentModel.Win32Exception: No connection could be made becaus
e the target machine actively refused
 it
```


Copyright 2016, ZeroC, Inc.

Note that, to successfully run client and server, the C# run time must be able to locate the `Ice.dll` library. (Under Windows, one way to ensure this is to copy the library into the current directory. Please consult the documentation for your C# run time to see how it locates libraries.)

See Also

- Client-Side Slice-to-C-Sharp Mapping
- Server-Side Slice-to-C-Sharp Mapping
- The `Ice.Application` Class
- The Current Object
- IceGrid

Copyright 2016, ZeroC, Inc.

Writing an Ice Application with Java

This page shows how to create an Ice application with Java.

On this page:

- Compiling a Slice Definition for Java
- Writing and Compiling a Server in Java
- Writing and Compiling a Client in Java
- Running Client and Server in Java

Compiling a Slice Definition for Java

The first step in creating our Java application is to compile our Slice definition to generate Java proxies and skeletons. You can compile the definition as follows:

```
$ mkdir generated
$ slice2java --output-dir generated Printer.ice
```

The `--output-dir` option instructs the compiler to place the generated files into the `generated` directory. This avoids cluttering the working directory with the generated files. The `slice2java` compiler produces a number of Java source files from this definition. The exact contents of these files do not concern us for now — they contain the generated code that corresponds to the `Printer` interface we defined in `Printer.ice`.

Writing and Compiling a Server in Java

To implement our `Printer` interface, we must create a servant class. By convention, a servant class uses the name of its interface with an `I`-suffix, so our servant class is called `PrinterI` and placed into a source file `PrinterI.java`:

Java

```
public class PrinterI extends Demo._PrinterDisp {
    public void
    printString(String s, Ice.Current current)
    {
        System.out.println(s);
    }
}
```

The `PrinterI` class inherits from a base class called `_PrinterDisp`, which is generated by the `slice2java` compiler. The base class is abstract and contains a `printString` method that accepts a string for the printer to print and a parameter of type `Ice.Current`. (For now we will ignore the `Ice.Current` parameter.) Our implementation of the `printString` method simply writes its argument to the terminal.

The remainder of the server code is in a source file called `Server.java`, shown in full here:

Copyright 2016, ZeroC, Inc.

Java

```
public class Server {
    public static void
    main(String[] args)
    {
        int status = 0;
        Ice.Communicator ic = null;
        try {
            ic = Ice.Util.initialize(args);
            Ice.ObjectAdapter adapter =
                ic.createObjectAdapterWithEndpoints("SimplePrinterAdapter", "defaul
t -p 10000");
            Ice.Object object = new PrinterI();
            adapter.add(object, ic.stringToIdentity("SimplePrinter"));
            adapter.activate();
            ic.waitForShutdown();
        } catch (Ice.LocalException e) {
            e.printStackTrace();
            status = 1;
        } catch (Exception e) {
            System.err.println(e.getMessage());
            status = 1;
        }
        if (ic != null) {
            // Clean up
            //
            try {
                ic.destroy();
            } catch (Exception e) {
                System.err.println(e.getMessage());
                status = 1;
            }
        }
        System.exit(status);
    }
}
```

Note the general structure of the code:

Copyright 2016, ZeroC, Inc.

Java

```
public class Server {
    public static void
    main(String[] args)
    {
        int status = 0;
        Ice.Communicator ic = null;
        try {

            // Server implementation here...

        } catch (Ice.LocalException e) {
            e.printStackTrace();
            status = 1;
        } catch (Exception e) {
            System.err.println(e.getMessage());
            status = 1;
        }
        if (ic != null) {
            // Clean up
            //
            try {
                ic.destroy();
            } catch (Exception e) {
                System.err.println(e.getMessage());
                status = 1;
            }
        }
        System.exit(status);
    }
}
```

The body of `main` contains a `try` block in which we place all the server code, followed by two `catch` blocks. The first block catches all exceptions that may be thrown by the Ice run time; the intent is that, if the code encounters an unexpected Ice run-time exception anywhere, the stack is unwound all the way back to `main`, which prints the exception and then returns failure to the operating system. The second block catches `Exception` exceptions; the intent is that, if we encounter a fatal error condition somewhere in our code, we can simply throw an exception with an error message. Again, this unwinds the stack all the way back to `main`, which prints the error message and then returns failure to the operating system.

Before the code exits, it destroys the communicator (if one was created successfully). Doing this is essential in order to correctly finalize the Ice run time: the program *must* call `destroy` on any communicator it has created; otherwise, undefined behavior results.

The body of our `try` block contains the actual server code:

Copyright 2016, ZeroC, Inc.

Java

```
            ic = Ice.Util.initialize(args);
            Ice.ObjectAdapter adapter =
                ic.createObjectAdapterWithEndpoints("SimplePrinterAdapter", "defaul
t -p 10000");
            Ice.Object object = new PrinterI();
            adapter.add(object, ic.stringToIdentity("SimplePrinter"));
            adapter.activate();
            ic.waitForShutdown();
```

The code goes through the following steps:

1. We initialize the Ice run time by calling `Ice.Util.initialize`. (We pass `args` to this call because the server may have command-line arguments that are of interest to the run time; for this example, the server does not require any command-line arguments.) The call to `initialize` returns an `Ice.Communicator` reference, which is the main object in the Ice run time.
2. We create an object adapter by calling `createObjectAdapterWithEndpoints` on the `Communicator` instance. The arguments we pass are `"SimplePrinterAdapter"` (which is the name of the adapter) and `"default -p 10000"`, which instructs the adapter to listen for incoming requests using the default protocol (TCP/IP) at port number 10000.
3. At this point, the server-side run time is initialized and we create a servant for our `Printer` interface by instantiating a `PrinterI` object.
4. We inform the object adapter of the presence of a new servant by calling `add` on the adapter; the arguments to `add` are the servant we have just instantiated, plus an identifier. In this case, the string `"SimplePrinter"` is the name of the Ice object. (If we had multiple printers, each would have a different name or, more correctly, a different *object identity*.)
5. Next, we activate the adapter by calling its `activate` method. (The adapter is initially created in a holding state; this is useful if we have many servants that share the same adapter and do not want requests to be processed until after all the servants have been instantiated.)
6. Finally, we call `waitForShutdown`. This call suspends the calling thread until the server implementation terminates, either by making a call to shut down the run time, or in response to a signal. (For now, we will simply interrupt the server on the command line when we no longer need it.)

Note that, even though there is quite a bit of code here, that code is essentially the same for all servers. You can put that code into a helper class and, thereafter, will not have to bother with it again. (Ice provides such a helper class, called `Ice.Application`.) As far as actual application code is concerned, the server contains only a few lines: seven lines for the definition of the `PrinterI` class, plus three lines to instantiate a `PrinterI` object and register it with the object adapter.

We can compile the server code as follows:

```
$ mkdir classes
$ javac -d classes -classpath classes:$ICE_HOME/lib/Ice.jar \
Server.java PrinterI.java generated/Demo/*.java
```

This compiles both our application code and the code that was generated by the Slice compiler. We assume that the `ICE_HOME` environment variable is set to the top-level directory containing the Ice run time. (For example, if you have installed Ice in `/opt/Ice`, set `ICE_HOME` to that path.) Note that Ice for Java uses the `ant` build environment to control building of source code. (`ant` is similar to `make`, but more flexible for Java applications.) You can have a look at the demo code that ships with Ice to see how to use this tool.

Writing and Compiling a Client in Java

The client code, in `Client.java`, looks very similar to the server. Here it is in full:

Copyright 2016, ZeroC, Inc.

Java

```
public class Client {
    public static void
    main(String[] args)
    {
        int status = 0;
        Ice.Communicator ic = null;
        try {
            ic = Ice.Util.initialize(args);
            Ice.ObjectPrx base = ic.stringToProxy("SimplePrinter:default -p 10000")
;
            Demo.PrinterPrx printer = Demo.PrinterPrxHelper.checkedCast(base);
            if (printer == null)
                throw new Error("Invalid proxy");

            printer.printString("Hello World!");
        } catch (Ice.LocalException e) {
            e.printStackTrace();
            status = 1;
        } catch (Exception e) {
            System.err.println(e.getMessage());
            status = 1;
        }
        if (ic != null) {
            // Clean up
            //
            try {
                ic.destroy();
            } catch (Exception e) {
                System.err.println(e.getMessage());
                status = 1;
            }
        }
        System.exit(status);
    }
}
```

Note that the overall code layout is the same as for the server: we use the same `try` and `catch` blocks to deal with errors. The code in the `try` block does the following:

1. As for the server, we initialize the Ice run time by calling `Ice.Util.initialize`.
2. The next step is to obtain a proxy for the remote printer. We create a proxy by calling `stringToProxy` on the communicator, with the string `"SimplePrinter:default -p 10000"`. Note that the string contains the object identity and the port number that were used by the server. (Obviously, hard-coding object identities and port numbers into our applications is a bad idea, but it will do for now; we will see more architecturally sound ways of doing this when we discuss IceGrid.)
3. The proxy returned by `stringToProxy` is of type `Ice.ObjectPrx`, which is at the root of the inheritance tree for interfaces and classes. But to actually talk to our printer, we need a proxy for a `Printer` interface, not an `Object` interface. To do this, we need to do a down-cast by calling `PrinterPrxHelper.checkedCast`. A checked cast sends a message to the server, effectively asking "is this a proxy for a `Printer` interface?" If so, the call returns a proxy of type `Demo::Printer`; otherwise, if the proxy denotes an interface of some other type, the call returns null.
4. We test that the down-cast succeeded and, if not, throw an error message that terminates the client.
5. We now have a live proxy in our address space and can call the `printString` method, passing it the time-honored `"Hello World!"` string. The server prints that string on its terminal.

Compiling the client looks much the same as for the server:

Copyright 2016, ZeroC, Inc.

```
$ javac -d classes -classpath classes:$ICE_HOME/lib/Ice.jar \
Client.java PrinterI.java generated/Demo/*.java
```

Running Client and Server in Java

To run client and server, we first start the server in a separate window:

```
$ java Server
```

At this point, we won't see anything because the server simply waits for a client to connect to it. We run the client in a different window:

```
$ java Client
$
```

The client runs and exits without producing any output; however, in the server window, we see the `"Hello World!"` that is produced by the printer. To get rid of the server, we interrupt it on the command line for now. (We will see cleaner ways to terminate a server in our discussion of `Ice.Application`.)

If anything goes wrong, the client will print an error message. For example, if we run the client without having first started the server, we get something like the following:

```
Ice.ConnectionRefusedException
       error = 0
    at
IceInternal.ConnectRequestHandler.getConnection(ConnectRequestHandler.java:240)
    at
IceInternal.ConnectRequestHandler.sendRequest(ConnectRequestHandler.java:138)
    at IceInternal.Outgoing.invoke(Outgoing.java:66)
    at Ice._ObjectDelM.ice_isA(_ObjectDelM.java:30)
    at Ice.ObjectPrxHelperBase.ice_isA(ObjectPrxHelperBase.java:111)
    at Ice.ObjectPrxHelperBase.ice_isA(ObjectPrxHelperBase.java:77)
    at Demo.HelloPrxHelper.checkedCast(HelloPrxHelper.java:228)
    at Client.run(Client.java:65)
Caused by: java.net.ConnectException: Connection refused
         ...
```

Note that, to successfully run client and server, your `CLASSPATH` must include the Ice library and the classes directory, for example:

```
$ export CLASSPATH=$CLASSPATH:./classes:$ICE_HOME/lib/Ice.jar
```

Please have a look at the demo applications that ship with Ice for the details for your platform.

See Also

- Client-Side Slice-to-Java Mapping
- Server-Side Slice-to-Java Mapping
- The `Ice.Application` Class
- The Current Object
- IceGrid

Copyright 2016, ZeroC, Inc.

Copyright 2016, ZeroC, Inc.

Writing an Ice Application with JavaScript

This page shows how to create an Ice client application with JavaScript.

On this page:

- Compiling a Slice Definition for JavaScript
- Writing a Client in JavaScript
- Running the Client in JavaScript

Compiling a Slice Definition for JavaScript

The first step in creating our JavaScript application is to compile our Slice definition to generate JavaScript proxies. You can compile the definition as follows:

```
$ slice2js Printer.ice
```

The `slice2js` compiler produces a single source file, `Printer.js`, from this definition. The exact contents of the source file do not concern us for now — it contains the generated code that corresponds to the `Printer` interface we defined in `Printer.ice`.

Writing a Client in JavaScript

The client code, in `Client.js`, is shown below in full:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var Ice = require("icejs").Ice;
var Demo = require("./Printer").Demo;

var ic;

Ice.Promise.try(
    function()
    {
        ic = Ice.initialize();
        var base = ic.stringToProxy("SimplePrinter:default -p 10000");
        return Demo.PrinterPrx.checkedCast(base).then(
            function(printer)
            {
                return printer.printString("Hello World!");
            });
    }
).finally(
    function()
    {
        if(ic)
        {
            return ic.destroy();
        }
    }
).exception(
    function(ex)
    {
        console.log(ex.toString());
        process.exit(1);
    });
```

The program begins with `require` statements that assign modules from the Ice run time and the generated code to convenient local variables. (These statements are necessary for use with NodeJS. Browser applications would omit these statements and load the modules a different way.)

The program begins with a call to `Ice.Promise.try` to launch a chain of promises (or futures) that handles the asynchronous nature of Ice invocations with a structure that resembles synchronous code.

1. The function passed to `try` is executed immediately. The body of this function begins by calling `Ice.initialize` to initialize the Ice run time. The call to `initialize` returns an `Ice.Communicator` reference, which is the main object in the Ice run time.
2. The next step is to obtain a proxy for the remote printer. We create a proxy by calling `stringToProxy` on the communicator, with the string `"SimplePrinter:default -p 10000"`. Note that the string contains the object identity and the port number that were used by the server. (Obviously, hard-coding object identities and port numbers into our applications is a bad idea, but it will do for now; we will see more architecturally sound ways of doing this when we discuss IceGrid.)
3. The proxy returned by `stringToProxy` is of type `Ice.ObjectPrx`, which is at the root of the inheritance tree for interfaces and classes. But to actually talk to our printer, we need a proxy for a `Demo::Printer` interface, not an `Object` interface. To do this, we need to do a down-cast by calling `Demo.PrinterPrx.checkedCast`. A checked cast sends a message to the server, effectively asking "is this a proxy for a `Demo::Printer` interface?" If so, the call returns a proxy of type `Demo::PrinterPrx`; otherwise, if the proxy denotes an interface of some other type, the call returns `null`.
4. The `checkedCast` function involves a remote invocation to the server, which means this function has asynchronous semantics and therefore it returns a new promise object.
5. We call `then` on the promise returned by `checkedCast` and supply a "success" function, meaning the code that's executed when `checkedCast` succeeds. This inner function accepts one argument, `printer`, representing a proxy to the newly-downcasted object, or `null` if the remote object doesn't support the `Printer` interface.

Copyright 2016, ZeroC, Inc.

6. Inside the success function, we now have a live proxy in our address space and can call the `printString` method, passing it the time-honored `"Hello World!"` string. The server prints that string on its terminal. Again, `printString` is a remote invocation, and it returns a promise that the success function passes along as its own return value.
7. The `then` function also returns a new promise which our outer function passes back to `try`. This outer promise is chained to the promise associated with the `printString` invocation; the outer promise completes successfully if and when the `printString` inv ocation completes successfully.
8. The function passed to `finally` is executed after the `try` block has completed, whether or not it completes successfully. If we created a communicator in the `try` block, we destroy it here. Doing this is essential in order to correctly finalize the Ice run time: the program *must* call `destroy` on any communicator it has created; otherwise, undefined behavior results. The `destroy` function has asynchronous semantics, so we return its promise to ensure no subsequent code is executed until `destroy` completes.
9. Lastly, the function passed to `exception` is the default exception handler for this entire promise chain.

Running the Client in JavaScript

The server must be started before the client. Since Ice for JavaScript does not currently include a complete server-side implementation, we need to use a server from another language mapping. In this case, we will use the C++ server:

```
$ server
```

At this point, we won't see anything because the server simply waits for a client to connect to it. We run the client in a different window:

```
$ node Client.js
$
```

The client runs and exits without producing any output; however, in the server window, we see the `"Hello World!"` that is produced by the printer. To get rid of the server, we interrupt it on the command line.

If anything goes wrong, the client will print an error message. For example, if we run the client without having first started the server, we get something like the following:

```
Ice::ConnectionRefusedException
    ice_cause: "Error: connect ECONNREFUSED"
    error: "ECONNREFUSED"
```

Note that, to successfully run the client, NodeJS must be able to locate the Ice for JavaScript modules. See the Ice for JavaScript installation instructions for more information.

See Also

- Client-Side Slice-to-Ruby Mapping
- IceGrid

Copyright 2016, ZeroC, Inc.

Writing an Ice Application with Objective-C

This page shows how to create an Ice application with Objective-C.

On this page:

- Compiling a Slice Definition for Objective-C
- Writing and Compiling a Server in Objective-C
- Writing and Compiling a Client in Objective-C
- Running Client and Server in Objective-C

Compiling a Slice Definition for Objective-C

The first step in creating our Objective-C application is to compile our Slice definition to generate Objective-C proxies and skeletons. You can compile the definition as follows:

```
$ slice2objc Printer.ice
```

The `slice2objc` compiler produces two Objective-C source files from this definition, `Printer.h` and `Printer.m`.

- `Printer.h`
  The `Printer.h` header file contains Objective-C type definitions that correspond to the Slice definitions for our `Printer` interface. This header file must be included in both the client and the server source code.
- `Printer.m`
  The `Printer.m` file contains the source code for our `Printer` interface. The generated source contains type-specific run-time support for both clients and servers. For example, it contains code that marshals parameter data (the string passed to the `printString` operation) on the client side and unmarshals that data on the server side.
  The `Printer.m` file must be compiled and linked into both client and server.

Writing and Compiling a Server in Objective-C

The source code for the server takes only a few lines and is shown in full here:

Objective-C

Copyright 2016, ZeroC, Inc.

```
#import <Ice/Ice.h>
#import <Printer.h>

#import <Foundation/NSAutoreleasePool.h>
#import <stdio.h>

@interface PrinterI : DemoPrinter <DemoPrinter>
@end

@implementation PrinterI
-(void) printString:(NSMutableString *)s
                    current:(ICECurrent *)current
{
    printf("%s\n", [s UTF8String]);
}
@end

int
main(int argc, char* argv[])
{
    NSAutoreleasePool * pool = [[NSAutoreleasePool alloc] init];

    int status = 1;
    id<ICECommunicator> communicator = nil;
    @try {
        communicator = [ICEUtil createCommunicator:&argc argv:argv];

        id<ICEObjectAdapter> adapter =
            [communicator createObjectAdapterWithEndpoints:
                                @"SimplePrinterAdapter"
                                endpoints:@"default -p 10000"];

        ICEObject *object = [[[PrinterI alloc] init] autorelease];
        [adapter add:object identity:[communicator
stringToIdentity:@"SimplePrinter"]];
        [adapter activate];

        [communicator waitForShutdown];

        status = 0;
    } @catch (NSException* ex) {
        NSLog(@"%@", ex);
    }

    @try {
        [communicator destroy];
    } @catch (NSException* ex) {
        NSLog(@"%@", ex);
    }

    [pool release];
    return status;
}
```

There appears to be a lot of code here for something as simple as a server that just prints a string. Do not be concerned by this: most of the

Copyright 2016, ZeroC, Inc.

preceding code is boiler plate that never changes. For this very simple server, the code is dominated by this boiler plate.

Every Ice source file starts with an include directive for `Ice.h`, which contains the definitions for the Ice run time. We also include `Printer.h`, which was generated by the Slice compiler and contains the Objective-C definitions for our printer interface. In addition, we import headers to allow us to use an autorelease pool and to produce output:

Objective-C

```
#import <Ice/Ice.h>
#import <Printer.h>

#import <Foundation/NSAutoreleasePool.h>
#import <stdio.h>
```

Our server implements a single printer servant, of type `PrinterI`. Looking at the generated code in `Printer.h`, we find the following (tidied up a little to get rid of irrelevant detail):

Objective-C

```
@protocol DemoPrinter <ICEObject>
-(void) printString:(NSMutableString *)s
                  current:(ICECurrent *)current;
@end

@interface DemoPrinter : ICEObject
// ...
@end
```

The `DemoPrinter` protocol and class definitions are generated by the Slice compiler. The protocol defines the `printString` method, which we must implement in our servant. The `DemoPrinter` class contains methods that are internal to the mapping, so we are not concerned with these. However, our servant must derive from this skeleton class:

Objective-C

```
@interface PrinterI : DemoPrinter <DemoPrinter>
@end

@implementation PrinterI
-(void) printString:(NSMutableString *)s
                  current:(ICECurrent *)current
{
    printf("%s\n", [s UTF8String]);
}
@end
```

As you can see, the implementation of the `printString` method is trivial: it simply writes its string argument to `stdout`.

Note that `printString` has a second parameter of type `ICECurrent`. The Ice run time passes additional information about an incoming request to the servant in this parameter. For now, we will ignore it.

What follows is the server main program. Note the general structure of the code:

Copyright 2016, ZeroC, Inc.

Objective-C

```
int
main(int argc, char* argv[])
{
    NSAutoreleasePool * pool = [[NSAutoreleasePool alloc] init];

    int status = 1;
    id<ICECommunicator> communicator = nil;
    @try {
        communicator = [ICEUtil createCommunicator:&argc argv:argv];

        // Server implementation here...

        status = 0;
    } @catch (NSException* ex) {
        NSLog(@"%@", ex);
    }

    @try {
        [communicator destroy];
    } @catch (NSException* ex) {
        NSLog(@"%@", ex);
    }

    [pool release];
    return status;
}
```

The body of `main` instantiates an autorelease pool, which it releases before returning to ensure that the program does not leak memory. `main` contains the declaration of two variables, `status` and `communicator`. The `status` variable contains the exit status of the program and the `communicator` variable, of type `id<ICECommunicator>`, contains the main handle to the Ice run time.

Following these declarations is a `try` block in which we place all the server code, followed by a `catch` handler that logs any unhandled exceptions.

Before returning, `main` executes a bit of cleanup code that calls the `destroy` method on the communicator. The cleanup call is outside the first `try` block for a reason: we must ensure that the Ice run time is finalized whether the code terminates normally or terminates due to an exception.

Failure to call `destroy` on the communicator before the program exits results in undefined behavior.

The body of the first `try` block contains the actual server code:

Copyright 2016, ZeroC, Inc.

Objective-C

```
        communicator = [ICEUtil createCommunicator:&argc argv:argv];

        id<ICEObjectAdapter> adapter =
            [communicator createObjectAdapterWithEndpoints:
                                  @"SimplePrinterAdapter"
                                  endpoints:@"default -p 10000"];

        ICEObject *object = [[[PrinterI alloc] init] autorelease];
        [adapter add:object identity:[communicator
stringToIdentity:@"SimplePrinter"]];
        [adapter activate];

        [communicator waitForShutdown];
```

The code goes through the following steps:

1. We initialize the Ice run time by calling `createCommunicator`. (We pass `argc` and `argv` to this call because the server may have command-line arguments that are of interest to the run time; for this example, the server does not require any command-line arguments.) The call to `createCommunicator` returns a value of type `id<ICECommunicator>`, which is the main object in the Ice run time.
2. We create an object adapter by calling `createObjectAdapterWithEndpoints` on the `Communicator` instance. The arguments we pass are `"SimplePrinterAdapter"` (which is the name of the adapter) and `"default -p 10000"`, which instructs the adapter to listen for incoming requests using the default protocol (TCP/IP) at port number 10000.
3. At this point, the server-side run time is initialized and we create a servant for our `Printer` interface by instantiating a `PrinterI` o bject.
4. We inform the object adapter of the presence of a new servant by calling `add` on the adapter; the arguments to `add` are the servant we have just instantiated, plus an identifier. In this case, the string `"SimplePrinter"` is the name of the Ice object. (If we had multiple printers, each would have a different name or, more correctly, a different *object identity*.)
5. Next, we activate the adapter by calling its `activate` method. (The adapter is initially created in a holding state; this is useful if we have many servants that share the same adapter and do not want requests to be processed until after all the servants have been instantiated.) The server starts to process incoming requests from clients as soon as the adapter is activated.
6. Finally, we call `waitForShutdown`. This call suspends the calling thread until the server implementation terminates, either by making a call to shut down the run time, or in response to a signal. (For now, we will simply interrupt the server on the command line when we no longer need it.)

Note that, even though there is quite a bit of code here, that code is essentially the same for all servers. You can put that code into a helper class and, thereafter, will not have to bother with it again. As far as actual application code is concerned, the server contains only a few lines: nine lines for the definition of the `PrinterI` class, plus three lines to instantiate a `PrinterI` object and register it with the object adapter.

Assuming that we have the server code in a file called `Server.m`, we can compile it as follows:

```
$ cc -c -I. -I$ICE_HOME/include Printer.m Server.m
```

This compiles both our application code and the code that was generated by the Slice compiler. We assume that the `ICE_HOME` environment variable is set to the top-level directory containing the Ice run time. (For example, if you have installed Ice in `/opt/Ice`, set `ICE_HOME` to that path.) Depending on your platform, you may have to add additional include directives or other options to the compiler; please see the demo programs that ship with Ice for the details.

Finally, we need to link the server into an executable:

```
$ c++ Printer.o Server.o -o server -L$ICE_HOME/lib -lIceObjC -framework Foundation
```

Again, depending on the platform, the actual list of libraries you need to link against may be longer. The demo programs that ship with Ice contain all the detail.

Copyright 2016, ZeroC, Inc.

Writing and Compiling a Client in Objective-C

The client code looks very similar to the server. Here it is in full:

Objective-C

```
#import <Ice/Ice.h>
#import <Printer.h>

#import <Foundation/NSAutoreleasePool.h>
#import <stdio.h>

int
main(int argc, char* argv[])
{
    NSAutoreleasePool * pool = [[NSAutoreleasePool alloc] init];

    int status = 1;
    id<ICECommunicator> communicator = nil;
    @try {
        communicator = [ICEUtil createCommunicator:&argc argv:argv];
        id<ICEObjectPrx> base = [communicator
stringToProxy:@"SimplePrinter:default -p 10000"];
        id<DemoPrinterPrx> printer = [DemoPrinterPrx checkedCast:base];
        if(!printer)
            [NSException raise:@"Invalid proxy" format:nil];

        [printer printString:@"Hello World!"];

        status = 0;
    } @catch (NSException* ex) {
        NSLog(@"%@", ex);
    }

    @try {
        [communicator destroy];
    } @catch (NSException* ex) {
        NSLog(@"%@", ex);
    }

    [pool release];
    return status;
}
```

Note that the overall code layout is the same as for the server: we include the headers for the Ice run time and the header generated by the Slice compiler, and we use the same `try` block and `catch` handlers to deal with errors.

The code in the `try` block does the following:

1. As for the server, we initialize the Ice run time by calling `createCommunicator`.
2. The next step is to obtain a proxy for the remote printer. We create a proxy by calling `stringToProxy` on the communicator, with the string `"SimplePrinter:default -p 10000"`. Note that the string contains the object identity and the port number that were used by the server. (Obviously, hard-coding object identities and port numbers into our applications is a bad idea, but it will do for now; we will see more architecturally sound ways of doing this when we discuss IceGrid.)
3. The proxy returned by `stringToProxy` is of type `id<ICEObjectPrx>`, which is at the root of the inheritance tree for interfaces

Copyright 2016, ZeroC, Inc.

and classes. But to actually talk to our printer, we need a proxy for a `Printer` interface, not an `Object` interface. To do this, we need to do a down-cast by calling the `checkedCast` class method on the `DemoPrinterPrx` class. A checked cast sends a message to the server, effectively asking "is this a proxy for a `Printer` interface?" If so, the call returns a proxy to a `Printer`; otherwise, if the proxy denotes an interface of some other type, the call returns a null proxy.

4. We test that the down-cast succeeded and, if not, throw an error message that terminates the client.
5. We now have a live proxy in our address space and can call the `printString` method, passing it the time-honored `"Hello World!"` string. The server prints that string on its terminal.

Compiling and linking the client looks much the same as for the server:

```
$ cc -c -I. -I$ICE_HOME/include Printer.m Client.m
$ c++ Printer.o Client.o -o client -L$ICE_HOME/lib -lIceObjC -framework Foundation
```

Running Client and Server in Objective-C

To run client and server, we first start the server in a separate window:

```
$ ./server
```

At this point, we won't see anything because the server simply waits for a client to connect to it. We run the client in a different window:

```
$ ./client
$
```

The client runs and exits without producing any output; however, in the server window, we see the `"Hello World!"` that is produced by the printer. To get rid of the server, we interrupt it on the command line for now.

If anything goes wrong, the client will print an error message. For example, if we run the client without having first started the server, we get:

```
Network.cpp:1218: Ice::ConnectionRefusedException:
connection refused: Connection refused
```

Note that, to successfully run client and server, you may have to set `DYLD_LIBRARY_PATH` to include the Ice library directory. Please see the installation instructions and the demo applications that ship with Ice Touch for details.

See Also

- Client-Side Slice-to-Objective-C Mapping
- Server-Side Slice-to-Objective-C Mapping
- The Current Object
- IceGrid

Copyright 2016, ZeroC, Inc.

Writing an Ice Application with PHP

This page shows how to create an Ice client application with PHP.

On this page:

- Compiling a Slice Definition for PHP
- Writing a Client in PHP
- Running the Client in PHP

Compiling a Slice Definition for PHP

The first step in creating our PHP application is to compile our Slice definition to generate PHP code. You can compile the definition as follows:

```
$ slice2php Printer.ice
```

The `slice2php` compiler produces a single source file, `Printer.php`, from this definition. The exact contents of the source file do not concern us for now — it contains the generated code that corresponds to the `Printer` interface we defined in `Printer.ice`.

Writing a Client in PHP

The client code, in `Client.php`, is shown below in full:

Copyright 2016, ZeroC, Inc.

PHP

```
<?php
require 'Ice.php';
require 'Printer.php';

$ic = null;
try
{
    $ic = Ice_initialize();
    $base = $ic->stringToProxy("SimplePrinter:default -p 10000");
    $printer = Demo_PrinterPrxHelper::checkedCast($base);
    if(!$printer)
        throw new RuntimeException("Invalid proxy");

    $printer->printString("Hello World!");
}
catch(Exception $ex)
{
    echo $ex;
}

if($ic)
{
    // Clean up
    try
    {
        $ic->destroy();
    }
    catch(Exception $ex)
    {
        echo $ex;
    }
}
?>
```

The program begins with `require` statements to load the Ice run-time definitions (`Ice.php`) and the code we generated from our Slice definition in the previous section (`Printer.php`).

The body of the main program contains a `try` block in which we place all the client code, followed by a `catch` block. The `catch` block catches all exceptions that may be thrown by the code; the intent is that, if the code encounters an unexpected run-time exception anywhere, the stack is unwound all the way back to the main program, which prints the exception and then returns failure to the operating system.

The body of our `try` block goes through the following steps:

1. We initialize the Ice run time by calling `Ice_initialize`. The call to `initialize` returns an `Ice_Communicator` reference, which is the main object in the Ice run time.
2. The next step is to obtain a proxy for the remote printer. We create a proxy by calling `stringToProxy` on the communicator, with the string `"SimplePrinter:default -p 10000"`. Note that the string contains the object identity and the port number that were used by the server. (Obviously, hard-coding object identities and port numbers into our applications is a bad idea, but it will do for now; we will see more architecturally sound ways of doing this when we discuss IceGrid.)
3. The proxy returned by `stringToProxy` is of type `Ice_ObjectPrx`, which is at the root of the inheritance tree for interfaces and classes. But to actually talk to our printer, we need a proxy for a `Demo::Printer` interface, not an `Object` interface. To do this, we need to do a down-cast by calling `Demo_PrinterPrxHelper::checkedCast`. A checked cast sends a message to the server, effectively asking "is this a proxy for a `Demo::Printer` interface?" If so, the call returns a proxy narrowed to the `Printer` interface ; otherwise, if the proxy denotes an interface of some other type, the call returns `null`.
4. We test that the down-cast succeeded and, if not, throw an exception that terminates the client.

Copyright 2016, ZeroC, Inc.

5. We now have a live proxy in our address space and can call the `printString` method, passing it the time-honored `"Hello World!"` string. The server prints that string on its terminal.

Before the code exits, it destroys the communicator (if one was created successfully). Doing this is essential in order to correctly finalize the Ice run time. If a script neglects to destroy the communicator, Ice destroys it automatically.

Running the Client in PHP

The server must be started before the client. Since Ice for PHP does not support server-side behavior, we need to use a server from another language mapping. In this case, we will use the C++ server:

```
$ server
```

At this point, we won't see anything because the server simply waits for a client to connect to it. We run the client in a different window using PHP's command-line interpreter:

```
$ php -f Client.php
$
```

The client runs and exits without producing any output; however, in the server window, we see the `"Hello World!"` that is produced by the printer. To get rid of the server, we interrupt it on the command line.

If anything goes wrong, the client will print an error message. For example, if we run the client without having first started the server, we get something like the following:

```
exception ::Ice::ConnectionRefusedException
{
    error = 111
}
```

Note that, to successfully run the client, the PHP interpreter must be able to locate the Ice extension for PHP. See the Ice for PHP installation instructions for more information.

See Also

- Client-Side Slice-to-PHP Mapping
- IceGrid

Copyright 2016, ZeroC, Inc.

Writing an Ice Application with Python

This page shows how to create an Ice application with Python.

On this page:

- Compiling a Slice Definition for Python
- Writing a Server in Python
- Writing a Client in Python
- Running Client and Server in Python

Compiling a Slice Definition for Python

The first step in creating our Python application is to compile our Slice definition to generate Python proxies and skeletons. You can compile the definition as follows:

```
$ slice2py Printer.ice
```

The `slice2py` compiler produces a single source file, `Printer_ice.py`, from this definition. The compiler also creates a Python package for the `Demo` module, resulting in a subdirectory named `Demo`. The exact contents of the source file do not concern us for now — it contains the generated code that corresponds to the `Printer` interface we defined in `Printer.ice`.

Writing a Server in Python

To implement our `Printer` interface, we must create a servant class. By convention, a servant class uses the name of its interface with an `I`-suffix, so our servant class is called `PrinterI`:

Python

```
class PrinterI(Demo.Printer):
    def printString(self, s, current=None):
        print s
```

The `PrinterI` class inherits from a base class called `Demo.Printer`, which is generated by the `slice2py` compiler. The base class is abstract and contains a `printString` method that accepts a string for the printer to print and a parameter of type `Ice.Current`. (For now we will ignore the `Ice.Current` parameter.) Our implementation of the `printString` method simply writes its argument to the terminal.

The remainder of the server code, in `Server.py`, follows our servant class and is shown in full here:

Copyright 2016, ZeroC, Inc.

Python

```
import sys, traceback, Ice
import Demo

class PrinterI(Demo.Printer):
    def printString(self, s, current=None):
        print s

status = 0
ic = None
try:
    ic = Ice.initialize(sys.argv)
    adapter = ic.createObjectAdapterWithEndpoints("SimplePrinterAdapter", "default
-p 10000")
    object = PrinterI()
    adapter.add(object, ic.stringToIdentity("SimplePrinter"))
    adapter.activate()
    ic.waitForShutdown()
except:
    traceback.print_exc()
    status = 1

if ic:
    # Clean up
    try:
        ic.destroy()
    except:
        traceback.print_exc()
        status = 1

sys.exit(status)
```

Note the general structure of the code:

Copyright 2016, ZeroC, Inc.

Python

```
status = 0
ic = None
try:

    # Server implementation here...

except:
    traceback.print_exc()
    status = 1

if ic:
    # Clean up
    try:
        ic.destroy()
    except:
        traceback.print_exc()
        status = 1

sys.exit(status)
```

The body of the main program contains a `try` block in which we place all the server code, followed by an `except` block. The `except` block catches all exceptions that may be thrown by the code; the intent is that, if the code encounters an unexpected run-time exception anywhere, the stack is unwound all the way back to the main program, which prints the exception and then returns failure to the operating system.

Before the code exits, it destroys the communicator (if one was created successfully). Doing this is essential in order to correctly finalize the Ice run time: the program *must* call `destroy` on any communicator it has created; otherwise, undefined behavior results.

The body of our `try` block contains the actual server code:

Python

```
ic = Ice.initialize(sys.argv)
adapter = ic.createObjectAdapterWithEndpoints("SimplePrinterAdapter", "default -p 1
0000")
object = PrinterI()
adapter.add(object, ic.stringToIdentity("SimplePrinter"))
adapter.activate()
ic.waitForShutdown()
```

The code goes through the following steps:

1. We initialize the Ice run time by calling `Ice.initialize`. (We pass `sys.argv` to this call because the server may have command-line arguments that are of interest to the run time; for this example, the server does not require any command-line arguments.) The call to `initialize` returns an `Ice.Communicator` reference, which is the main object in the Ice run time.
2. We create an object adapter by calling `createObjectAdapterWithEndpoints` on the `Communicator` instance. The arguments we pass are `"SimplePrinterAdapter"` (which is the name of the adapter) and `"default -p 10000"`, which instructs the adapter to listen for incoming requests using the default protocol (TCP/IP) at port number 10000.
3. At this point, the server-side run time is initialized and we create a servant for our `Printer` interface by instantiating a `PrinterI` object.
4. We inform the object adapter of the presence of a new servant by calling `add` on the adapter; the arguments to `add` are the servant we have just instantiated, plus an identifier. In this case, the string `"SimplePrinter"` is the name of the Ice object. (If we had multiple printers, each would have a different name or, more correctly, a different *object identity*.)
5. Next, we activate the adapter by calling its `activate` method. (The adapter is initially created in a holding state; this is useful if we

Copyright 2016, ZeroC, Inc.

have many servants that share the same adapter and do not want requests to be processed until after all the servants have been instantiated.)
6. Finally, we call `waitForShutdown`. This call suspends the calling thread until the server implementation terminates, either by making a call to shut down the run time, or in response to a signal. (For now, we will simply interrupt the server on the command line when we no longer need it.)

Note that, even though there is quite a bit of code here, that code is essentially the same for all servers. You can put that code into a helper class and, thereafter, will not have to bother with it again. (Ice provides such a helper class, called `Ice.Application`.) As far as actual application code is concerned, the server contains only a few lines: three lines for the definition of the `PrinterI` class, plus two lines to instantiate a `PrinterI` object and register it with the object adapter.

Writing a Client in Python

The client code, in `Client.py`, looks very similar to the server. Here it is in full:

Python

```
import sys, traceback, Ice
import Demo

status = 0
ic = None
try:
    ic = Ice.initialize(sys.argv)
    base = ic.stringToProxy("SimplePrinter:default -p 10000")
    printer = Demo.PrinterPrx.checkedCast(base)
    if not printer:
        raise RuntimeError("Invalid proxy")

    printer.printString("Hello World!")
except:
    traceback.print_exc()
    status = 1

if ic:
    # Clean up
    try:
        ic.destroy()
    except:
        traceback.print_exc()
        status = 1

sys.exit(status)
```

Note that the overall code layout is the same as for the server: we use the same `try` and `except` blocks to deal with errors. The code in the `try` block does the following:

1. As for the server, we initialize the Ice run time by calling `Ice.initialize`.
2. The next step is to obtain a proxy for the remote printer. We create a proxy by calling `stringToProxy` on the communicator, with the string `"SimplePrinter:default -p 10000"`. Note that the string contains the object identity and the port number that were used by the server. (Obviously, hard-coding object identities and port numbers into our applications is a bad idea, but it will do for now; we will see more architecturally sound ways of doing this when we discuss IceGrid.)
3. The proxy returned by `stringToProxy` is of type `Ice.ObjectPrx`, which is at the root of the inheritance tree for interfaces and classes. But to actually talk to our printer, we need a proxy for a `Demo::Printer` interface, not an `Object` interface. To do this, we need to do a down-cast by calling `Demo.PrinterPrx.checkedCast`. A checked cast sends a message to the server, effectively asking "is this a proxy for a `Demo::Printer` interface?" If so, the call returns a proxy of type `Demo.PrinterPrx`; otherwise, if the proxy denotes an interface of some other type, the call returns `None`.
4. We test that the down-cast succeeded and, if not, throw an error message that terminates the client.

Copyright 2016, ZeroC, Inc.

5. We now have a live proxy in our address space and can call the `printString` method, passing it the time-honored `"Hello World!"` string. The server prints that string on its terminal.

Running Client and Server in Python

To run client and server, we first start the server in a separate window:

```
$ python Server.py
```

At this point, we won't see anything because the server simply waits for a client to connect to it. We run the client in a different window:

```
$ python Client.py
$
```

The client runs and exits without producing any output; however, in the server window, we see the `"Hello World!"` that is produced by the printer. To get rid of the server, we interrupt it on the command line for now. (We will see cleaner ways to terminate a server in our discussion of `Ice.Application`.)

If anything goes wrong, the client will print an error message. For example, if we run the client without having first started the server, we get something like the following:

```
Traceback (most recent call last):
  File "Client.py", line 10, in ?
    printer = Demo.PrinterPrx.checkedCast(base)
  File "Printer_ice.py", line 43, in checkedCast
    return Demo.PrinterPrx.ice_checkedCast(proxy, '::Demo::Printer', facet)
ConnectionRefusedException: Ice.ConnectionRefusedException:
Connection refused
```

Note that, to successfully run the client and server, the Python interpreter must be able to locate the Ice extension for Python. See the Ice for Python installation instructions for more information.

See Also

- Client-Side Slice-to-Python Mapping
- Server-Side Slice-to-Python Mapping
- The `Ice.Application` Class
- The Current Object
- IceGrid

Copyright 2016, ZeroC, Inc.

Writing an Ice Application with Ruby

This page shows how to create an Ice client application with Ruby.

On this page:

- Compiling a Slice Definition for Ruby
- Writing a Client in Ruby
- Running the Client in Ruby

Compiling a Slice Definition for Ruby

The first step in creating our Ruby application is to compile our Slice definition to generate Ruby proxies. You can compile the definition as follows:

```
$ slice2rb Printer.ice
```

The `slice2rb` compiler produces a single source file, `Printer.rb`, from this definition. The exact contents of the source file do not concern us for now — it contains the generated code that corresponds to the `Printer` interface we defined in `Printer.ice`.

Writing a Client in Ruby

The client code, in `Client.rb`, is shown below in full:

Copyright 2016, ZeroC, Inc.

Ruby

```
require 'Printer.rb'

status = 0
ic = nil
begin
    ic = Ice::initialize(ARGV)
    base = ic.stringToProxy("SimplePrinter:default -p 10000")
    printer = Demo::PrinterPrx::checkedCast(base)
    if not printer
        raise "Invalid proxy"
    end

    printer.printString("Hello World!")
rescue
    puts $!
    puts $!.backtrace.join("\n")
    status = 1
end

if ic
    # Clean up
    begin
        ic.destroy()
    rescue
        puts $!
        puts $!.backtrace.join("\n")
        status = 1
    end
end

exit(status)
```

The program begins with a `require` statement, which loads the Ruby code we generated from our Slice definition in the previous section. It is not necessary for the client to explicitly load the `Ice` module because `Printer.rb` does that for you.

The body of the main program contains a `begin` block in which we place all the client code, followed by a `rescue` block. The `rescue` block catches all exceptions that may be thrown by the code; the intent is that, if the code encounters an unexpected run-time exception anywhere, the stack is unwound all the way back to the main program, which prints the exception and then returns failure to the operating system.

The body of our `begin` block goes through the following steps:

1. We initialize the Ice run time by calling `Ice::initialize`. (We pass `ARGV` to this call because the client may have command-line arguments that are of interest to the run time; for this example, the client does not require any command-line arguments.) The call to `initialize` returns an `Ice::Communicator` reference, which is the main object in the Ice run time.
2. The next step is to obtain a proxy for the remote printer. We create a proxy by calling `stringToProxy` on the communicator, with the string `"SimplePrinter:default -p 10000"`. Note that the string contains the object identity and the port number that were used by the server. (Obviously, hard-coding object identities and port numbers into our applications is a bad idea, but it will do for now; we will see more architecturally sound ways of doing this when we discuss IceGrid.)
3. The proxy returned by `stringToProxy` is of type `Ice::ObjectPrx`, which is at the root of the inheritance tree for interfaces and classes. But to actually talk to our printer, we need a proxy for a `Demo::Printer` interface, not an `Object` interface. To do this, we need to do a down-cast by calling `Demo::PrinterPrx::checkedCast`. A checked cast sends a message to the server, effectively asking "is this a proxy for a `Demo::Printer` interface?" If so, the call returns a proxy of type `Demo::PrinterPrx`; otherwise, if the proxy denotes an interface of some other type, the call returns `nil`.
4. We test that the down-cast succeeded and, if not, throw an error message that terminates the client.
5. We now have a live proxy in our address space and can call the `printString` method, passing it the time-honored `"Hello`

Copyright 2016, ZeroC, Inc.

`World!"` string. The server prints that string on its terminal.

Before the code exits, it destroys the communicator (if one was created successfully). Doing this is essential in order to correctly finalize the Ice run time: the program *must* call `destroy` on any communicator it has created; otherwise, undefined behavior results.

Running the Client in Ruby

The server must be started before the client. Since Ice for Ruby does not support server-side behavior, we need to use a server from another language mapping. In this case, we will use the C++ server:

```
$ server
```

At this point, we won't see anything because the server simply waits for a client to connect to it. We run the client in a different window:

```
$ ruby Client.rb
$
```

The client runs and exits without producing any output; however, in the server window, we see the `"Hello World!"` that is produced by the printer. To get rid of the server, we interrupt it on the command line.

If anything goes wrong, the client will print an error message. For example, if we run the client without having first started the server, we get something like the following:

```
exception ::Ice::ConnectionRefusedException
{
    error = 111
}
```

Note that, to successfully run the client, the Ruby interpreter must be able to locate the Ice extension for Ruby. See the Ice for Ruby installation instructions for more information.

See Also

- Client-Side Slice-to-Ruby Mapping
- IceGrid

Copyright 2016, ZeroC, Inc.

Writing an Ice Application with Visual Basic

This page shows how to create an Ice application with Visual Basic.

On this page:

- Visual Basic Development Process
- Compiling a Slice Definition for Visual Basic
- Writing and Compiling a Server in Visual Basic
- Writing and Compiling a Client in Visual Basic
- Running Client and Server in Visual Basic

Visual Basic Development Process

As of version 3.3, Ice no longer includes a separate compiler to create Visual Basic source code from Slice definitions. Instead, you need to use the Slice-to-C# compiler `slice2cs` to create C# source code and compile the generated C# source code with a C# compiler into a DLL that contains the compiled generated code for your Slice definitions. Your Visual Basic application then links with this DLL and the Ice for .NET DLL (`Ice.dll`).

This approach works not only with Visual Basic, but with any language that targets the .NET run time. However, ZeroC does not provide support for languages other than C# and Visual Basic.

The following illustration demonstrates this development process:

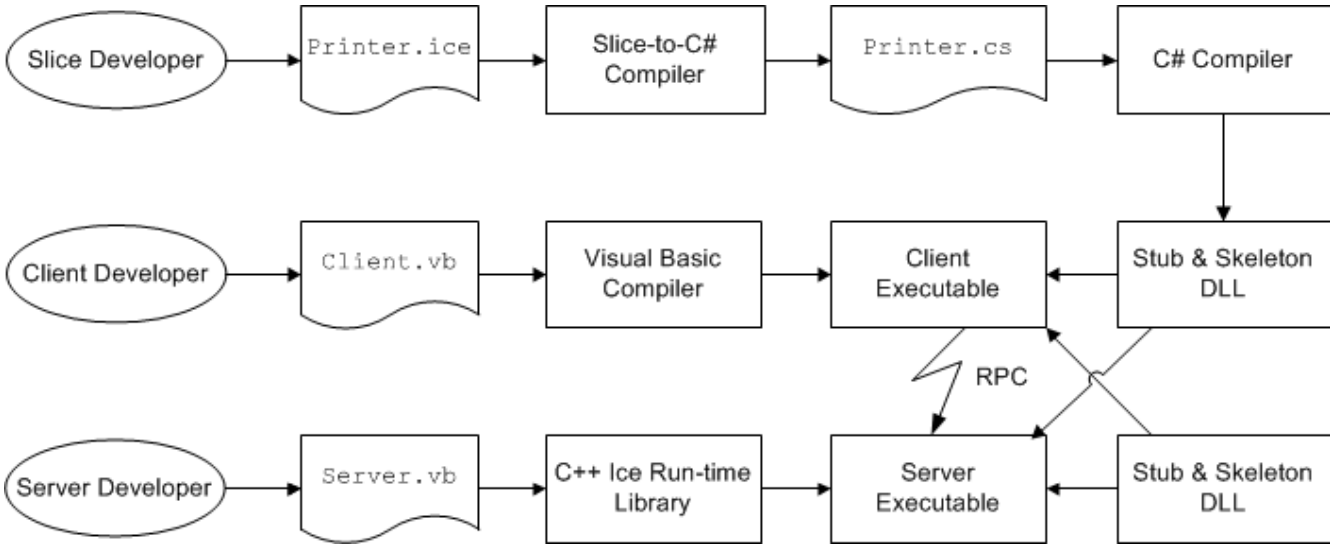

Developing a Visual Basic application with Ice.

Compiling a Slice Definition for Visual Basic

The first step in creating our VB application is to compile our Slice definition to generate proxies and skeletons. You can compile the definition as follows:

```
> mkdir generated
> slice2cs --output-dir generated Printer.ice
```

The `--output-dir` option instructs the compiler to place the generated files into the `generated` directory. This avoids cluttering the working directory with the generated files. The `slice2cs` compiler produces a single source file, `Printer.cs`, from this definition. The exact contents of this file do not concern us for now — it contains the generated code that corresponds to the `Printer` interface we defined in `Printer.ice`.

We now need to compile this generated code into a DLL:

Copyright 2016, ZeroC, Inc.

```
> csc /reference:Ice.dll /lib:%ICE_HOME%\bin /t:library /out:Printer.dll
generated\Printer.cs
```

This creates a DLL called `Printer.dll` that contains the code we generated from the Slice definitions.

Writing and Compiling a Server in Visual Basic

To implement our `Printer` interface, we must create a servant class. By convention, a servant class uses the name of its interface with an `I`-suffix, so our servant class is called `PrinterI` and placed into a source file `Server.vb`:

Visual Basic

```
Imports System
Imports Demo

Public Class PrinterI
    Inherits PrinterDisp_

    Public Overloads Overrides Sub printString( _
                                ByVal s As String, _
                                ByVal current As Ice.Current)
        Console.WriteLine(s)
    End Sub

End Class
```

The `PrinterI` class inherits from a base class called `PrinterDisp_`, which is generated by the `slice2cs` compiler. The base class is abstract and contains a `printString` method that accepts a string for the printer to print and a parameter of type `Ice.Current`. (For now we will ignore the `Ice.Current` parameter.) Our implementation of the `printString` method simply writes its argument to the terminal.

The remainder of the server code follows in `Server.vb` and is shown in full here:

Copyright 2016, ZeroC, Inc.

Visual Basic

```
Module Server

    Public Sub Main(ByVal args() As String)

        Dim status As Integer = 0
        Dim ic As Ice.Communicator = Nothing
        Try
            ic = Ice.Util.initialize(args)
            Dim adapter As Ice.ObjectAdapter = _
                ic.createObjectAdapterWithEndpoints("SimplePrinterAdapter", "defaul
t -p 10000")
            Dim obj As Ice.Object = New PrinterI
            adapter.add(obj, ic.stringToIdentity("SimplePrinter"))
            adapter.activate()
            ic.waitForShutdown()
        Catch e As Exception
            Console.Error.WriteLine(e)
            status = 1
        End Try
        If Not ic Is Nothing Then
            ' Clean up
            '
            Try
                ic.destroy()
            Catch e As Exception
                Console.Error.WriteLine(e)
                status = 1
            End Try
        End If
        Environment.Exit(status)
    End Sub

End module
```

Note the general structure of the code:

Copyright 2016, ZeroC, Inc.

Visual Basic

```
Module Server

    Public Sub Main(ByVal args() As String)

        Dim status As Integer = 0
        Dim ic As Ice.Communicator = Nothing
        Try

            ' Server implementation here...

        Catch e As Exception
            Console.Error.WriteLine(e)
            status = 1
        End Try
        If Not ic Is Nothing Then
            ' Clean up
            '
            Try
                ic.destroy()
            Catch e As Exception
                Console.Error.WriteLine(e)
                status = 1
            End Try
        End If
        Environment.Exit(status)
    End Sub

End module
```

The body of `Main` contains a `Try` block in which we place all the server code, followed by a `Catch` block. The catch block catches all exceptions that may be thrown by the code; the intent is that, if the code encounters an unexpected run-time exception anywhere, the stack is unwound all the way back to `Main`, which prints the exception and then returns failure to the operating system.

Before the code exits, it destroys the communicator (if one was created successfully). Doing this is essential in order to correctly finalize the Ice run time: the program *must* call `destroy` on any communicator it has created; otherwise, undefined behavior results.

The body of our `Try` block contains the actual server code:

Visual Basic

```
            ic = Ice.Util.initialize(args)
            Dim adapter As Ice.ObjectAdapter = _
                ic.createObjectAdapterWithEndpoints("SimplePrinterAdapter", "defaul
t -p 10000")
            Dim obj As Ice.Object = New PrinterI
            adapter.add(obj, ic.stringToIdentity("SimplePrinter"))
            adapter.activate()
            ic.waitForShutdown()
```

The code goes through the following steps:

Copyright 2016, ZeroC, Inc.

1. We initialize the Ice run time by calling `Ice.Util.initialize`. (We pass `args` to this call because the server may have command-line arguments that are of interest to the run time; for this example, the server does not require any command-line arguments.) The call to `initialize` returns an `Ice::Communicator` reference, which is the main object in the Ice run time.
2. We create an object adapter by calling `createObjectAdapterWithEndpoints` on the `Communicator` instance. The arguments we pass are `"SimplePrinterAdapter"` (which is the name of the adapter) and `"default -p 10000"`, which instructs the adapter to listen for incoming requests using the default protocol (TCP/IP) at port number 10000.
3. At this point, the server-side run time is initialized and we create a servant for our `Printer` interface by instantiating a `PrinterI` object.
4. We inform the object adapter of the presence of a new servant by calling `add` on the adapter; the arguments to `add` are the servant we have just instantiated, plus an identifier. In this case, the string `"SimplePrinter"` is the name of the servant. (If we had multiple printers, each would have a different name or, more correctly, a different *object identity*.)
5. Next, we activate the adapter by calling its `activate` method. (The adapter is initially created in a holding state; this is useful if we have many servants that share the same adapter and do not want requests to be processed until after all the servants have been instantiated.)
6. Finally, we call `waitForShutdown`. This call suspends the calling thread until the server implementation terminates, either by making a call to shut down the run time, or in response to a signal. (For now, we will simply interrupt the server on the command line when we no longer need it.)

Note that, even though there is quite a bit of code here, that code is essentially the same for all servers. You can put that code into a helper class and, thereafter, will not have to bother with it again. (Ice provides such a helper class, called `Ice.Application`.) As far as actual application code is concerned, the server contains only a few lines: ten lines for the definition of the `PrinterI` class, plus three lines to instantiate a `PrinterI` object and register it with the object adapter.

We can compile the server code as follows:

```
> vbc /reference:Ice.dll /libpath:%ICE_HOME%\bin /reference:Printer.dll
/out:server.exe Server.vb
```

This compiles our application code and links it with the Ice run time and the DLL we generated earlier. We assume that the `ICE_HOME` environment variable is set to the top-level directory containing the Ice run time. (For example, if you have installed Ice in `C:\Ice`, set `ICE_HOME` to that path.)

Writing and Compiling a Client in Visual Basic

The client code, in `Client.vb`, looks very similar to the server. Here it is in full:

Copyright 2016, ZeroC, Inc.

Visual Basic

```
Imports System
Imports Demo

Module Client

    Public Sub Main(ByVal args() As String)
        Dim status As Integer = 0
        Dim ic As Ice.Communicator = Nothing
        Try
            ic = Ice.Util.initialize(args)
            Dim obj As Ice.ObjectPrx = ic.stringToProxy("SimplePrinter:default -p 1
0000")
            Dim printer As PrinterPrx = PrinterPrxHelper.checkedCast(obj)
            If printer Is Nothing Then
                Throw New ApplicationException("Invalid proxy")
            End If

            printer.printString("Hello World!")
        Catch e As Exception
            Console.Error.WriteLine(e)
            status = 1
        End Try
        If Not ic Is Nothing Then
            ' Clean up
            '
            Try
                ic.destroy()
            Catch e As Exception
                Console.Error.WriteLine(e)
                status = 1
            End Try
        End If
        Environment.Exit(status)
    End Sub

End Module
```

Note that the overall code layout is the same as for the server: we use the same `Try` and `Catch` blocks to deal with errors. The code in the `Try` block does the following:

1. As for the server, we initialize the Ice run time by calling `Ice.Util.initialize`.
2. The next step is to obtain a proxy for the remote printer. We create a proxy by calling `stringToProxy` on the communicator, with the string `"SimplePrinter:default -p 10000"`. Note that the string contains the object identity and the port number that were used by the server. (Obviously, hard-coding object identities and port numbers into our applications is a bad idea, but it will do for now; we will see more architecturally sound ways of doing this when we discuss IceGrid.)
3. The proxy returned by `stringToProxy` is of type `Ice.ObjectPrx`, which is at the root of the inheritance tree for interfaces and classes. But to actually talk to our printer, we need a proxy for a `Printer` interface, not an `Object` interface. To do this, we need to do a down-cast by calling `PrinterPrxHelper.checkedCast`. A checked cast sends a message to the server, effectively asking "is this a proxy for a `Printer` interface?" If so, the call returns a proxy of type `Demo::Printer`; otherwise, if the proxy denotes an interface of some other type, the call returns null.
4. We test that the down-cast succeeded and, if not, throw an error message that terminates the client.
5. We now have a live proxy in our address space and can call the `printString` method, passing it the time-honored `"Hello World!"` string. The server prints that string on its terminal.

Copyright 2016, ZeroC, Inc.

Compiling the client looks much the same as for the server:

```
> vbc /reference:Ice.dll /libpath:%ICE_HOME%\bin /reference:Printer.dll
/out:client.exe Client.vb
```

Running Client and Server in Visual Basic

To run client and server, we first start the server in a separate window:

```
> server.exe
```

At this point, we won't see anything because the server simply waits for a client to connect to it. We run the client in a different window:

```
> client.exe
>
```

The client runs and exits without producing any output; however, in the server window, we see the `"Hello World!"` that is produced by the printer. To get rid of the server, we interrupt it on the command line for now. (We will see cleaner ways to terminate a server in our discussion of `Ice.Application`.)

If anything goes wrong, the client will print an error message. For example, if we run the client without having first started the server, we get something like the following:

```
Ice.ConnectionRefusedException
    error = 0
   at IceInternal.ProxyFactory.checkRetryAfterException(LocalException ex, Referenc
e ref, Int32 cnt)
   at Ice.ObjectPrxHelperBase.handleException__(ObjectDel_ delegate, LocalException
 ex, Int32 cnt)
   at Ice.ObjectPrxHelperBase.ice_isA(String id__, Dictionary`2 context__, Boolean
explicitContext__)
   at Ice.ObjectPrxHelperBase.ice_isA(String id__)
   at Demo.PrinterPrxHelper.checkedCast(ObjectPrx b)
   at Client.Main(String[] args)
Caused by: System.ComponentModel.Win32Exception: No connection could be made becaus
e the target machine actively refused
 it
```

Note that, to successfully run client and server, the VB run time must be able to locate the `Ice.dll` library. (Under Windows, one way to ensure this is to copy the library into the current directory. Please consult the documentation for your VB run time to see how it locates libraries.)

See Also

- Client-Side Slice-to-C-Sharp Mapping
- Server-Side Slice-to-C-Sharp Mapping
- The `Ice.Application` Class
- The Current Object
- IceGrid

Copyright 2016, ZeroC, Inc.

The Slice Language

Here, we present the Slice language. Slice (Specification Language for Ice) is the fundamental abstraction mechanism for separating object interfaces from their implementations. Slice establishes a contract between client and server that describes the types and object interfaces used by an application. This description is independent of the implementation language, so it does not matter whether the client is written in the same language as the server.

Even though Slice is an acronym, it is pronounced as a single syllable, like a slice of bread.

Slice definitions are compiled for a particular implementation language by a compiler. The compiler translates the language-independent definitions into language-specific type definitions and APIs. These types and APIs are used by the developer to provide application functionality and to interact with Ice. The translation algorithms for various implementation languages are known as *language mappings*. Currently, Ice defines language mappings for C++, Java, C#, Python, Objective-C, Ruby, and PHP.

Because Slice describes interfaces and types (but not implementations), it is a purely declarative language; there is no way to write executable statements in Slice.

Slice definitions focus on object interfaces, the operations supported by those interfaces, and exceptions that may be raised by operations. In addition, Slice offers features for object persistence. This requires quite a bit of supporting machinery; in particular, much of Slice is concerned with the definition of data types. This is because data can be exchanged between client and server only if their types are defined in Slice. You cannot exchange arbitrary C++ data between client and server because it would destroy the language independence of Ice. However, you can always create a Slice type definition that corresponds to the C++ data you want to send, and then you can transmit the Slice type.

We present the full syntax and semantics of Slice here. Because much of Slice is based on C++ and Java, we focus on those areas where Slice differs from C++ or Java or constrains the equivalent C++ or Java feature in some way. Slice features that are identical to C++ and Java are mentioned mostly by example.

Topics

- Slice Compilation
- Slice Source Files
- Lexical Rules
- Modules
- Basic Types
- User-Defined Types
- Constants and Literals
- Interfaces, Operations, and Exceptions
- Classes
- Forward Declarations
- Optional Data Members
- Type IDs
- Operations on Object
- Local Types
- Names and Scoping
- Metadata
- Serializable Objects
- Deprecating Slice Definitions
- Using the Slice Compilers
- Slice Checksums
- Generating Slice Documentation
- Slice Keywords
- Slice Metadata Directives
- Slice for a Simple File System

Copyright 2016, ZeroC, Inc.

Slice Compilation

On this page:

- Compilation
- Single Development Environment for Client and Server
- Different Development Environments for Client and Server

Compilation

A Slice compiler produces source files that must be combined with application code to produce client and server executables.

The outcome of the development process is a client executable and a server executable. These executables can be deployed anywhere, whether the target environments use the same or different operating systems and whether the executables are implemented using the same or different languages. The only constraint is that the host machines must provide the necessary run-time environment, such as any required dynamic libraries, and that connectivity can be established between them.

Single Development Environment for Client and Server

The figure below shows the situation when both client and server are developed in C++. The Slice compiler generates two files from a Slice definition in a source file `Printer.ice`: a header file (`Printer.h`) and a source file (`Printer.cpp`)

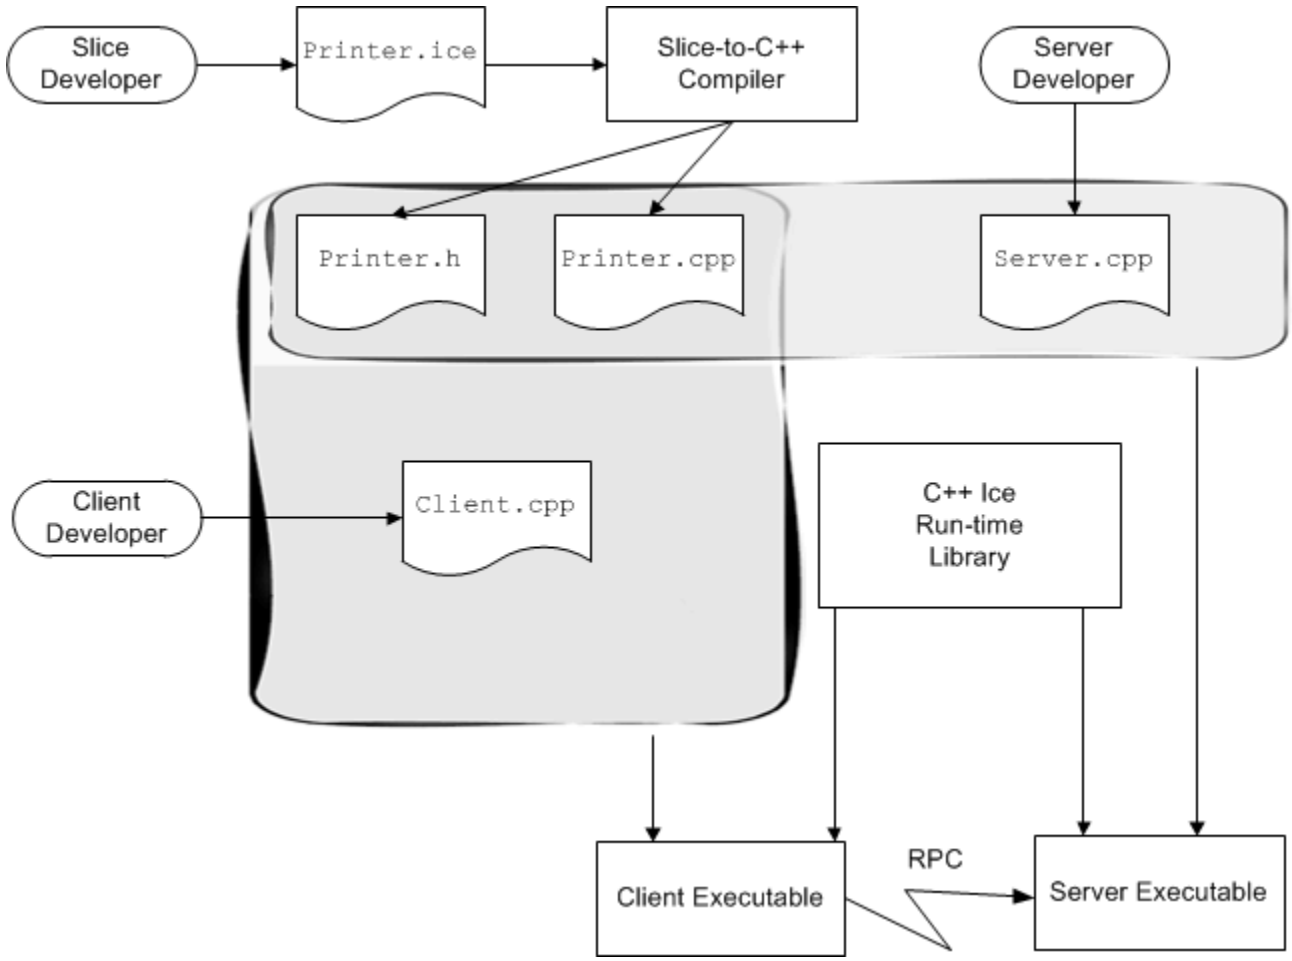

.

Development process if client and server share the same development environment.

- The `Printer.h` header file contains definitions that correspond to the types used in the Slice definition. It is included in the source code of both client and server to ensure that client and server agree about the types and interfaces used by the application.
- The `Printer.cpp` source file provides an API to the client for sending messages to remote objects. The client source code (`Client.cpp`, written by the client developer) contains the client-side application logic. The generated source code and the client code are compiled and linked into the client executable.

Copyright 2016, ZeroC, Inc.

The `Printer.cpp` source file also contains source code that provides an up-call interface from the Ice run time into the server code written by the developer and provides the connection between the networking layer of Ice and the application code. The server implementation file (`Server.cpp`, written by the server developer) contains the server-side application logic (the object implementations, properly termed *servants*). The generated source code and the implementation source code are compiled and linked into the server executable.

Both client and server also link with an Ice library that provides the necessary run-time support.

You are not limited to a single implementation of a client or server. For example, you can build multiple servers, each of which implements the same interfaces but uses different implementations (for example, with different performance characteristics). Multiple such server implementations can coexist in the same system. This arrangement provides one fundamental scalability mechanism in Ice: if you find that a server process starts to bog down as the number of objects increases, you can run an additional server for the same interfaces on a different machine. Such *federated* servers provide a single logical service that is distributed over a number of processes on different machines. Each server in the federation implements the same interfaces but hosts different object instances. (Of course, federated servers must somehow ensure consistency of any databases they share across the federation.)

Ice also provides support for *replicated* servers. Replication permits multiple servers to each implement the same set of object instances. This improves performance and scalability (because client load can be shared over a number of servers) as well as redundancy (because each object is implemented in more than one server).

Different Development Environments for Client and Server

Client and server cannot share any source or binary components if they are developed in different languages. For example, a client written in Java cannot include a C++ header file.

This figure shows the situation when a client written in Java and the corresponding server is written in C++. In this case, the client and server developers are completely independent, and each uses his or her own development environment and language mapping. The only link between client and server developers is the Slice definition each one uses.

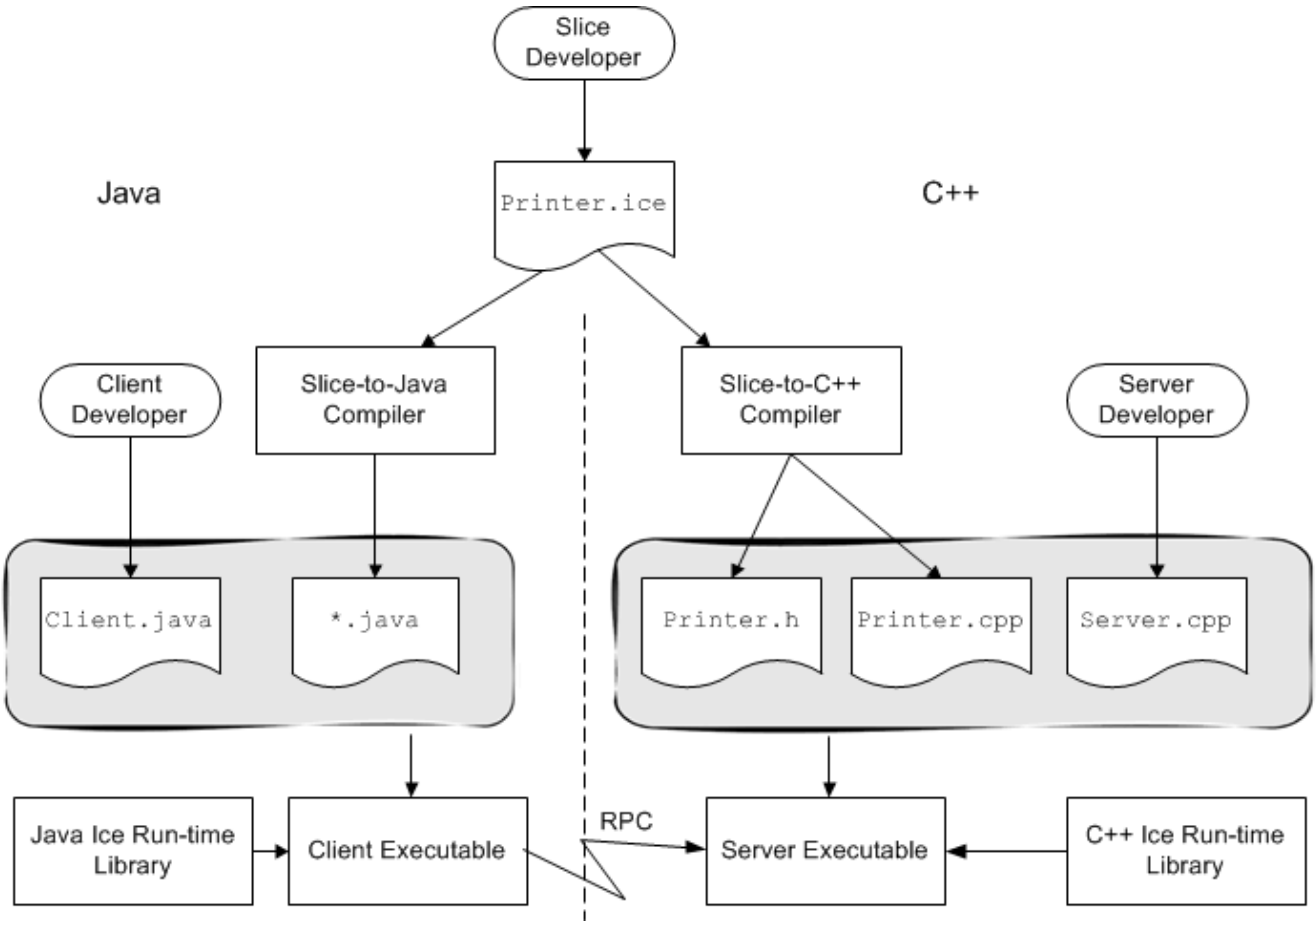

Development process for different development environments.

For Java, the slice compiler creates a number of files whose names depend on the names of various Slice constructs. (These files are collectively referred to as `*.java` in the above figure.)

Copyright 2016, ZeroC, Inc.

See Also

- Using the Slice Compilers

Copyright 2016, ZeroC, Inc.

Slice Source Files

Slice defines a number of rules for the naming and contents of Slice source files.

On this page:

- File Naming
- File Format
- Preprocessing
  - Detecting Ice Versions
  - Detecting Slice Compilers
- Definition Order

File Naming

Files containing Slice definitions must end in a `.ice` file extension, for example, `Clock.ice` is a valid file name. Other file extensions are rejected by the compilers.

For case-insensitive file systems (such as DOS), the file extension may be written as uppercase or lowercase, so `Clock.ICE` is legal. For case-sensitive file systems (such as Unix), `Clock.ICE` is illegal. (The extension must be in lowercase.)

File Format

Slice is a free-form language so you can use spaces, horizontal and vertical tab stops, form feeds, and newline characters to lay out your code in any way you wish. (White space characters are token separators). Slice does not attach semantics to the layout of a definition. You may wish to follow the style we have used for the Slice examples throughout this book.

Slice files can be ASCII text files or use the UTF-8 character encoding with a byte order marker (BOM) at the beginning of each file. However, Slice identifiers are limited to ASCII letters and digits; non-ASCII letters can appear only in comments and string literals.

Preprocessing

Slice supports the same preprocessor directives as C++, so you can use directives such as `#include` and macro definitions. However, Slice permits `#include` directives only at the beginning of a file, before any Slice definitions.

If you use `#include` directives, it is a good idea to protect them with guards to prevent double inclusion of a file:

Slice

```
// File Clock.ice
#ifndef _CLOCK_ICE
#define _CLOCK_ICE

// #include directives here...
// Definitions here...

#endif _CLOCK_ICE
```

The following `#pragma` directive offers a simpler way to achieve the same result:

Copyright 2016, ZeroC, Inc.

Slice

```
// File Clock.ice
#pragma once

// #include directives here...
// Definitions here...
```

`#include` directives permit a Slice definition to use types defined in a different source file. The Slice compilers parse all of the code in a source file, including the code in subordinate `#include` files. However, the compilers generate code only for the top-level file(s) nominated on the command line. You must separately compile subordinate `#include` files to obtain generated code for all the files that make up your Slice definition.

Note that you should avoid `#include` with double quotes:

Slice

```
#include "Clock.ice" // Not recommended!
```

While double quotes will work, the directory in which the preprocessor tries to locate the file can vary depending on the operating system, so the included file may not always be found where you expect it. Instead, use angle brackets (<>); you can control which directories are searched for the file with the `-I` option of the Slice compiler.

Also note that, if you include a path separator in a `#include` directive, you must use a forward slash:

Slice

```
#include <SliceDefs/Clock.ice>  // OK
```

You cannot use a backslash in `#include` directives:

Slice

```
#include <SliceDefs\Clock.ice>  // Illegal
```

Detecting Ice Versions

As of Ice 3.5, the Slice compilers define the preprocessor macro `__ICE_VERSION__` with a numeric representation of the Ice version. The value of this macro is the same as the C++ macro `ICE_INT_VERSION`. You can use this macro to make your Slice definitions backward-compatible with older Ice releases, while still taking advantage of newer Ice features when possible. For example, the Slice definition shown below makes use of custom enumerator values:

Copyright 2016, ZeroC, Inc.

Slice

```
#if defined(__ICE_VERSION__) && __ICE_VERSION__ >= 030500
enum Fruit { Apple, Pear = 3, Orange };
#else
enum Fruit { Apple, Pear, Orange };
#endif
```

Although this example is intended to show how to use the `ICE_VERSION` macro, it also highlights a potential pitfall that you must be aware of when trying to maintain backward compatibility: the two definitions of `Fruit` are not wire-compatible.

Detecting Slice Compilers

As of Ice 3.5, each Slice compiler defines its own macro so that you can customize your Slice code for certain language mappings. The following macros are defined by their respective compilers:

```
__SLICE2JAVA__
__SLICE2JS__
__SLICE2CPP__
__SLICE2CS__
__SLICE2PY__
__SLICE2PHP__
__SLICE2RB__
__SLICE2FREEZE__
__SLICE2FREEZEJ__
__SLICE2HTML__
__TRANSFORMDB__
__DUMPDB__
```

For example, .NET developers may elect to avoid the use of default values for structure members because the presence of default values changes the C# mapping of the structure from `struct` to `class`:

Slice

```
struct Record {
    // ...
#if __SLICE2CS__
    bool active;
#else
    bool active = true;
#endif
};
```

Definition Order

Slice constructs, such as modules, interfaces, or type definitions, can appear in any order you prefer. However, identifiers must be declared before they can be used.

See Also

- Using the Slice Compilers

Copyright 2016, ZeroC, Inc.

Lexical Rules

Slice's lexical rules are very similar to those of C++ and Java, except for some differences for identifiers.

On this page:

- Comments
- Keywords
- Identifiers
  - Case Sensitivity
  - Identifiers That Are Keywords
  - Escaped Identifiers
  - Reserved Identifiers

Comments

Slice definitions permit both the C and the C++ style of writing comments:

Slice

```
/*
 * C-style comment.
 */

// C++-style comment extending to the end of this line.
```

Keywords

Slice uses a number of keywords, which must be spelled in lowercase. For example, `class` and `dictionary` are keywords and must be spelled as shown. There are two exceptions to this lowercase rule: `Object` and `LocalObject` are keywords and must be capitalized as shown.

Identifiers

Identifiers begin with an alphabetic character followed by any number of alphabetic characters or digits. Underscores are also permitted in identifiers with the following limitations:

- an identifier cannot begin or end with an underscore
- an identifier cannot contain multiple consecutive underscores

Given these rules, the identifier `get_account_name` is legal but not `_account`, `account_`, or `get__account`.

Slice identifiers are restricted to the ASCII range of alphabetic characters and cannot contain non-English letters, such as Å. (Supporting non-ASCII identifiers would make it very difficult to map Slice to target languages that lack support for this feature.)

Case Sensitivity

Identifiers are case-insensitive but must be capitalized consistently. For example, `TimeOfDay` and `TIMEOFDAY` are considered the same identifier within a naming scope. However, Slice enforces consistent capitalization. After you have introduced an identifier, you must capitalize it consistently throughout; otherwise, the compiler will reject it as illegal. This rule exists to permit mappings of Slice to languages that ignore case in identifiers as well as to languages that treat differently capitalized identifiers as distinct.

Identifiers That Are Keywords

You can define Slice identifiers that are keywords in one or more implementation languages. For example, `switch` is a perfectly good Slice identifier but is a C++ and Java keyword. Each language mapping defines rules for dealing with such identifiers. The solution typically involves using a prefix to map away from the keyword. For example, the Slice identifier `switch` is mapped to `_cpp_switch` in C++ and `_switch` in Java.

Copyright 2016, ZeroC, Inc.

The rules for dealing with keywords can result in hard-to-read source code. Identifiers such as `native`, `throw`, or `export` will clash with C++ or Java keywords (or both). To make life easier for yourself and others, try to avoid Slice identifiers that are implementation language keywords. Keep in mind that mappings for new languages may be added to Ice in the future. While it is not reasonable to expect you to compile a list of all keywords in all popular programming languages, you should make an attempt to avoid at least common keywords. Slice identifiers such as `self`, `import`, and `while` are definitely not a good idea.

Escaped Identifiers

It is possible to use a Slice keyword as an identifier by prefixing the keyword with a backslash, for example:

Slice

```
struct dictionary {     // Error!
    // ...
};

struct \dictionary {    // OK
    // ...
};

struct \foo {           // Legal, same as "struct foo"
    // ...
};
```

The backslash escapes the usual meaning of a keyword; in the preceding example, `\dictionary` is treated as the identifier `dictionary`. The escape mechanism exists to permit keywords to be added to the Slice language over time with minimal disruption to existing specifications: if a pre-existing specification happens to use a newly-introduced keyword, that specification can be fixed by simply prepending a backslash to the new keyword. Note that, as a matter of style, you should avoid using Slice keywords as identifiers (even though the backslash escapes allow you to do this).

It is legal (though redundant) to precede an identifier that is not a keyword with a backslash — the backslash is ignored in that case.

Reserved Identifiers

Slice reserves the identifier `Ice` and all identifiers beginning with `Ice` (in any capitalization) for the Ice implementation. For example, if you try to define a type named `Icecream`, the Slice compiler will issue an error message.

You can suppress this behavior by using the `--ice` compiler option, which enables definition of identifiers beginning with `Ice`. However, do not use this option unless you are compiling the Slice definitions for the Ice run time itself.

Slice identifiers ending in any of the suffixes `Helper`, `Holder`, `Prx`, and `Ptr` are also reserved. These endings are used by the various language mappings and are reserved to prevent name clashes in the generated code.

See Also

- Slice Keywords

Copyright 2016, ZeroC, Inc.

Modules

On this page:

- Modules Reduce Clutter
- Modules are Mandatory
- Reopening Modules
- Module Mapping
- The Ice Module

Modules Reduce Clutter

A common problem in large systems is pollution of the global namespace: over time, as isolated systems are integrated, name clashes become quite likely. Slice provides the `module` construct to alleviate this problem:

Slice

```
module ZeroC {
    module Client {
        // Definitions here...
    };
    module Server {
        // Definitions here...
    };
};
```

A module can contain any legal Slice construct, including other module definitions. Using modules to group related definitions together avoids polluting the global namespace and makes accidental name clashes quite unlikely. (You can use a well-known name, such as a company or product name, as the name of the outermost module.)

Modules are Mandatory

Slice requires all definitions to be nested inside a module, that is, you cannot define anything other than a module at global scope. For example, the following is illegal:

Slice

```
interface I {   // Error: only modules can appear at global scope
    // ...
};
```

Definitions at global scope are prohibited because they cause problems with some implementation languages (such as Python, which does not have a true global scope).

Throughout the Ice manual, you will occasionally see Slice definitions that are not nested inside a module. This is to keep the examples short and free of clutter. Whenever you see such a definition, assume that it is nested in a module.

Reopening Modules

Modules can be reopened:

Copyright 2016, ZeroC, Inc.

Slice

```
module ZeroC {
    // Definitions here...
};

// Possibly in a different source file:

module ZeroC {  // OK, reopened module
    // More definitions here...
};
```

Reopened modules are useful for larger projects: they allow you to split the contents of a module over several different source files. The advantage of doing this is that, when a developer makes a change to one part of the module, only files dependent on the changed part need be recompiled (instead of having to recompile all files that use the module).

Module Mapping

Modules map to a corresponding scoping construct in each programming language. (For example, for C++ and C#, Slice modules map to namespaces whereas, for Java, they map to packages.) This allows you to use an appropriate C++ `using` or Java `import` declaration to avoid excessively long identifiers in your source code.

The Ice Module

APIs for the Ice run time, apart from a small number of language-specific calls that cannot be expressed in Ice, are defined in the `Ice` module. In other words, most of the Ice API is actually expressed as Slice definitions. The advantage of doing this is that a single Slice definition is sufficient to define the API for the Ice run time for all supported languages. The respective language mapping rules then determine the exact shape of each Ice API for each implementation language.

We will incrementally explore the contents of the `Ice` module throughout this manual.

See Also

- Slice Source Files

Copyright 2016, ZeroC, Inc.

Basic Types

On this page:

- Built-In Basic Types
- Integer Types
- Floating-Point Types
- Strings
- Booleans
- Bytes

Built-In Basic Types

Slice provides a number of built-in basic types, as shown in this table:

Type	Range of Mapped Type	Size of Mapped Type
`bool`	`false` or `true`	1bit
`byte`	-128-127 or 0-255 [a]	8 bits
`short`	-2^{15} to $2^{15}-1$	16 bits
`int`	-2^{31} to $2^{31}-1$	32 bits
`long`	-2^{63} to $2^{63}-1$	64 bits
`float`	IEEE single-precision	32 bits
`double`	IEEE double-precision	64 bits
`string`	All Unicode characters, excluding the character with all bits zero.	Variable-length

[a] The range depends on whether `byte` maps to a signed or an unsigned type.

All the basic types (except `byte`) are subject to changes in representation as they are transmitted between clients and servers. For example, a `long` value is byte-swapped when sent from a little-endian to a big-endian machine. Similarly, strings undergo translation in representation if they are sent, for example, from an EBCDIC to an ASCII implementation, and the characters of a string may also change in size. (Not all architectures use 8-bit characters). However, these changes are transparent to the programmer and do exactly what is required.

Integer Types

Slice provides integer types `short`, `int`, and `long`, with 16-bit, 32-bit, and 64-bit ranges, respectively. Note that, on some architectures, any of these types may be mapped to a native type that is wider. Also note that no unsigned types are provided. (This choice was made because unsigned types are difficult to map into languages without native unsigned types, such as Java. In addition, the unsigned integers add little value to a language. (See [1] for a good treatment of the topic.)

Floating-Point Types

These types follow the IEEE specification for single- and double-precision floating-point representation [2]. If an implementation cannot support IEEE format floating-point values, the Ice run time converts values into the native floating-point representation (possibly at a loss of precision or even magnitude, depending on the capabilities of the native floating-point format).

Strings

Slice strings use the Unicode character set. The only character that cannot appear inside a string is the zero character.

This decision was made as a concession to C++, with which it becomes impossibly difficult to manipulate strings with embedded zero characters using standard library routines, such as `strlen` or `strcat`.

The Slice data model does not have the concept of a null string (in the sense of a C++ null pointer). This decision was made because null strings are difficult to map to languages without direct support for this concept (such as Python). Do not design interfaces that depend on a null string to indicate "not there" semantics. If you need the notion of an optional string, use a class, a sequence of strings, or use an empty

Copyright 2016, ZeroC, Inc.

string to represent the idea of a null string. (Of course, the latter assumes that the empty string is not otherwise used as a legitimate string value by your application.)

Booleans

Boolean values can have only the values `false` and `true`. Language mappings use the corresponding native boolean type if one is available.

Bytes

The Slice type `byte` is an (at least) 8-bit type that is guaranteed not to undergo any changes in representation as it is transmitted between address spaces. This guarantee permits exchange of binary data such that it is not tampered with in transit. All other Slice types are subject to changes in representation during transmission.

See Also

- Sequences
- Classes

References

1. Lakos, J. 1996. Large-Scale C++ Software Design. Reading, MA: Addison-Wesley.
2. Institute of Electrical and Electronics Engineers. 1985. *IEEE 754-1985 Standard for Binary Floating-Point Arithmetic*. Piscataway, NJ: Institute of Electrical and Electronic Engineers.

Copyright 2016, ZeroC, Inc.

User-Defined Types

In addition to providing the built-in basic types, Slice allows you to define complex types: enumerations, structures, sequences, and dictionaries.

Topics

- Enumerations
- Structures
- Sequences
- Dictionaries

Copyright 2016, ZeroC, Inc.

Enumerations

Enumeration Syntax and Semantics

A Slice enumerated type definition looks identical to C++:

Slice

```
enum Fruit { Apple, Pear, Orange };
```

This definition introduces a type named `Fruit` that becomes a new type in its own right. Slice guarantees that the values of enumerators increase from left to right, so `Apple` compares less than `Pear` in every language mapping. By default, the first enumerator has a value of zero, with sequentially increasing values for subsequent enumerators.

Slice enumerator symbols enter the enclosing namespace scope, so the following is illegal:

Slice

```
enum Fruit { Apple, Pear, Orange };
enum ComputerBrands { Apple, IBM, Sun, HP };    // Apple redefined
```

The example below shows how to refer to an enumerator from a different scope:

Slice

```
module M {
    enum Color { red, green, blue };
};
module N {
    struct Pixel {
        M::Color c = M::blue;
    };
};
```

Slice does not permit empty enumerations.

Custom Enumerator Values

Slice also permits you to assign custom values to enumerators:

Slice

```
const int PearValue = 7;
enum Fruit { Apple = 0, Pear = PearValue, Orange };
```

Custom values must be unique and non-negative, and may refer to Slice constants of integer types. If no custom value is specified for an enumerator, its value is one greater than the enumerator that immediately precedes it. In the example above, `Orange` has the value 8.

The maximum value for an enumerator value is the same as the maximum value for `int`, $2^{31} - 1$.

Copyright 2016, ZeroC, Inc.

Slice does not require custom enumerator values to be declared in increasing order:

Slice

```
enum Fruit { Apple = 5, Pear = 3, Orange = 1 };   // Legal
```

Note however that when there is an inconsistency between the declaration order and the numerical order of the enumerators, the behavior of comparison operations may vary between language mappings.

For an application that is still using version 1.0 of the Ice encoding, changing the definition of an enumerated type may break backward compatibility with existing applications. For more information, please refer to the encoding rules for enumerated types.

See Also

- Structures
- Sequences
- Dictionaries
- Constants and Literals

Copyright 2016, ZeroC, Inc.

Structures

Slice supports structures containing one or more named members of arbitrary type, including user-defined complex types. For example:

Slice

```
struct TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
};
```

As in C++, this definition introduces a new type called `TimeOfDay`. Structure definitions form a namespace, so the names of the structure members need to be unique only within their enclosing structure.

Data member definitions using a named type are the only construct that can appear inside a structure. It is impossible to, for example, define a structure inside a structure:

Slice

```
struct TwoPoints {
    struct Point {      // Illegal!
                short x;
                short y;
    };
    Point       coord1;
    Point       coord2;
};
```

This rule applies to Slice in general: type definitions cannot be nested (except for modules, which do support nesting). The reason for this rule is that nested type definitions can be difficult to implement for some target languages and, even if implementable, greatly complicate the scope resolution rules. For a specification language, such as Slice, nested type definitions are unnecessary – you can always write the above definitions as follows (which is stylistically cleaner as well):

Slice

```
struct Point {
    short x;
    short y;
};

struct TwoPoints {      // Legal (and cleaner!)
    Point coord1;
    Point coord2;
};
```

You can specify a default value for a data member that has one of the following types:

- An integral type (`byte`, `short`, `int`, `long`)

Copyright 2016, ZeroC, Inc.

- A floating point type (`float` or `double`)
- `string`
- `bool`
- `enum`

For example:

Slice

```
struct Location {
    string name;
    Point pt;
    bool display = true;
    string source = "GPS";
};
```

The legal syntax for literal values is the same as for Slice constants, and you may also use a constant as a default value. The language mapping guarantees that data members are initialized to their declared default values using a language-specific mechanism.

See Also

- Modules
- Basic Types
- Enumerations
- Sequences
- Dictionaries
- Constants and Literals

Copyright 2016, ZeroC, Inc.

Sequences

On this page:

- Sequence Syntax and Semantics
- Using Sequences for Optional Values

Sequence Syntax and Semantics

Sequences are variable-length collections of elements:

Slice

```
sequence<Fruit> FruitPlatter;
```

A sequence can be empty—that is, it can contain no elements, or it can hold any number of elements up to the memory limits of your platform.

Sequences can contain elements that are themselves sequences. This arrangement allows you to create lists of lists:

Slice

```
sequence<FruitPlatter> FruitBanquet;
```

Sequences are used to model a variety of collections, such as vectors, lists, queues, sets, bags, or trees. (It is up to the application to decide whether or not order is important; by discarding order, a sequence serves as a set or bag.)

Using Sequences for Optional Values

Using a sequence to model an optional value is unnecessary with the introduction of optional data members and optional parameters in Ice 3.5.

One particular use of sequences has become idiomatic, namely, the use of a sequence to indicate an optional value. For example, we might have a `Part` structure that records the details of the parts that go into a car. The structure could record things such as the name of the part, a description, weight, price, and other details. Spare parts commonly have a serial number, which we can model as a `long` value. However, some parts, such as simple screws, often do not have a serial number, so what are we supposed to put into the serial number field of a screw? There are a number of options for dealing with this situation:

- **Use a sentinel value, such as zero, to indicate the "no serial number" condition.**
  This approach is workable, provided that a sentinel value is actually available. While it may seem unlikely that anyone would use a serial number of zero for a part, it is not impossible. And, for other values, such as a temperature value, all values in the range of their type can be legal, so no sentinel value is available.
- **Change the type of the serial number from `long` to `string`.**
  Strings come with their own built-in sentinel value, namely the empty string, so we can use an empty string to indicate the "no serial number" case. This is workable but not ideal: we should not have to change the natural data type of something to `string` just so we get a sentinel value.
- **Add an indicator as to whether the contents of the serial number are valid:**

Copyright 2016, ZeroC, Inc.

Slice

```
struct Part {
    string name;
    string description;
    // ...
    bool   serialIsValid;  // true if part has serial number
    long   serialNumber;
};
```

This is guaranteed to get you into trouble eventually: sooner or later, some programmer will forget to check whether the serial number is valid before using it and create havoc.

- **Use a sequence to model the optional field.**
  This technique uses the following convention:

Slice

```
sequence<long> SerialOpt;

struct Part {
    string    name;
    string    description;
    // ...
    SerialOpt serialNumber; // optional: zero or one element
};
```

By convention, the `Opt` suffix is used to indicate that the sequence is used to model an optional value. If the sequence is empty, the value is obviously not there; if it contains a single element, that element is the value. The obvious drawback of this scheme is that someone could put more than one element into the sequence. This could be rectified by adding a special-purpose Slice construct for optional values. However, optional values are not used frequently enough to justify the complexity of adding a dedicated language feature. (As we will see in Classes, you can also use class hierarchies to model optional fields.)

See Also

- Enumerations
- Structures
- Dictionaries
- Constants and Literals
- Classes

Copyright 2016, ZeroC, Inc.

Dictionaries

On this page:

- Dictionary Syntax and Semantics
- Allowable Types for Dictionary Keys and Values

Dictionary Syntax and Semantics

A dictionary is a mapping from a key type to a value type.

For example:

Slice

```
struct Employee {
    long   number;
    string firstName;
    string lastName;
};

dictionary<long, Employee> EmployeeMap;
```

This definition creates a dictionary named `EmployeeMap` that maps from an employee number to a structure containing the details for an employee. Whether or not the key type (the employee number, of type `long` in this example) is also part of the value type (the `Employee` structure in this example) is up to you — as far as Slice is concerned, there is no need to include the key as part of the value.

Dictionaries can be used to implement sparse arrays, or any lookup data structure with non-integral key type. Even though a sequence of structures containing key-value pairs could be used to model the same thing, a dictionary is more appropriate:

- A dictionary clearly signals the intent of the designer, namely, to provide a mapping from a domain of values to a range of values. (A sequence of structures of key-value pairs does not signal that same intent as clearly.)
- At the programming language level, sequences are implemented as vectors (or possibly lists), that is, they are not well suited to model sparsely populated domains and require a linear search to locate an element with a particular value. On the other hand, dictionaries are implemented as a data structure (typically a hash table or red-black tree) that supports efficient searching in $O(\log n)$ average time or better.

Allowable Types for Dictionary Keys and Values

The key type of a dictionary need not be an integral type. For example, we could use the following definition to translate the names of the days of the week:

Slice

```
dictionary<string, string> WeekdaysEnglishToGerman;
```

The server implementation would take care of initializing this map with the key-value pairs `Monday-Montag`, `Tuesday-Dienstag`, and so on.

The value type of a dictionary can be any Slice type. However, the key type of a dictionary is limited to one of the following types:

- Integral types (`byte`, `short`, `int`, `long`, `bool`)
- `string`
- `enum`
- Structures containing only data members of legal key types

Copyright 2016, ZeroC, Inc.

Other complex types, such as dictionaries, and floating-point types (`float` and `double`) cannot be used as the key type. Complex types are disallowed because they complicate the language mappings for dictionaries, and floating-point types are disallowed because representational changes of values as they cross machine boundaries can lead to ill-defined semantics for equality.

See Also

- Basic Types
- Enumerations
- Structures
- Sequences
- Constants and Literals

Copyright 2016, ZeroC, Inc.

Constants and Literals

On this page:

- Allowable Types for Constants
- Boolean constants
- Integer literals
- Floating-point literals
- String literals
- Constant Expressions

Allowable Types for Constants

Slice allows you to define constants for the following types:

- An integral type (`bool`, `byte`, `short`, `int`, `long`)
- A floating point type (`float` or `double`)
- `string`
- `enum`

Here are a few examples:

Slice

```
const bool      AppendByDefault = true;
const byte      LowerNibble = 0x0f;
const string    Advice = "Don't Panic!";
const short     TheAnswer = 42;
const double    PI = 3.1416;

enum Fruit { Apple, Pear, Orange };
const Fruit     FavoriteFruit = Pear;
```

The syntax for literals is the same as for C++ and Java (with a few minor exceptions).

Boolean constants

Boolean constants can only be initialized with the keywords `false` and `true`. (You cannot use `0` and `1` to represent `false` and `true`.)

Integer literals

Integer literals can be specified in decimal, octal, or hexadecimal notation.

For example:

Slice

```
const byte TheAnswer = 42;
const byte TheAnswerInOctal = 052;
const byte TheAnswerInHex = 0x2A;       // or 0x2a
```

Be aware that, if you interpret `byte` as a number instead of a bit pattern, you may get different results in different languages. For example, for C++, `byte` maps to `unsigned char` whereas, for Java, `byte` maps to `byte`, which is a signed type.

Note that suffixes to indicate long and unsigned constants (`l`, `L`, `u`, `U`, used by C++) are illegal:

Copyright 2016, ZeroC, Inc.

Slice

```
const long Wrong = 0u;          // Syntax error
const long WrongToo = 1000000L; // Syntax error
```

The value of an integer literal must be within the range of its constant type, as shown in the Built-In Basic Types table; otherwise the compiler will issue a diagnostic.

Floating-point literals

Floating-point literals use C++ syntax, except that you cannot use an `l` or `L` suffix to indicate an extended floating-point constant; however, `f` and `F` are legal (but are ignored).

Here are a few examples:

Slice

```
const float P1 = -3.14f;    // Integer & fraction, with suffix
const float P2 = +3.1e-3;   // Integer, fraction, and exponent
const float P3 = .1;        // Fraction part only
const float P4 = 1.;        // Integer part only
const float P5 = .9E5;      // Fraction part and exponent
const float P6 = 5e2;       // Integer part and exponent
```

Floating-point literals must be within the range of the constant type (`float` or `double`); otherwise, the compiler will issue a diagnostic.

String literals

String literals support the same escape sequences as C++.

Here are some examples:

Copyright 2016, ZeroC, Inc.

Slice

```
const string AnOrdinaryString = "Hello World!";

const string DoubleQuote =      "\"";
const string TwoSingleQuotes =  "'\'";     // ' and \' are OK
const string Newline =          "\n";
const string CarriageReturn =   "\r";
const string HorizontalTab =    "\t";
const string VerticalTab =      "\v";
const string FormFeed =         "\f";
const string Alert =            "\a";
const string Backspace =        "\b";
const string QuestionMark =     "\?";
const string Backslash =        "\\";

const string OctalEscape =    "\007";     // Same as \a
const string HexEscape =      "\x07";     // Ditto
const string Universal1 =     "\u0041";   // Same as A
const string Universal2 = "\U00000041";   // Ditto
```

Note that Slice has no concept of a null string:

Slice

```
const string NullString = 0;    // Illegal!
```

Null strings simply do not exist in Slice and, therefore, do not exist as a legal value for a string anywhere in the Ice platform. The reason for this decision is that null strings do not exist in many programming languages.

Many languages other than C and C++ use a byte array as the internal string representation. Null strings do not exist (and would be very difficult to map) in such languages.

Constant Expressions

A constant definition may also refer to another constant. It is not necessary for both constants to have the same Slice type, but the value of the existing constant must be compatible with the type of the constant being defined.

Consider the examples below:

Slice

```
const int SIZE = 500;

const int DEFAULT_SIZE = SIZE; // OK
const short SHORT_SIZE = SIZE; // OK
const byte BYTE_SIZE = SIZE;   // ERROR
```

The `DEFAULT_SIZE` constant is legal because it has the same type as `SIZE`, and `SHORT_SIZE` is legal because the value of `SIZE` (`500`) is within the range of the Slice `short` type. However, `BYTE_SIZE` is illegal because the value of `SIZE` is outside the range of the `byte` type.

Copyright 2016, ZeroC, Inc.

See Also

- Enumerations
- Structures
- Sequences
- Dictionaries

Copyright 2016, ZeroC, Inc.

Interfaces, Operations, and Exceptions

The central focus of Slice is on defining interfaces, for example:

Slice

```
struct TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
};

interface Clock {
    TimeOfDay getTime();
    void setTime(TimeOfDay time);
};
```

This definition defines an interface type called `Clock`. The interface supports two operations: `getTime` and `setTime`. Clients access an object supporting the `Clock` interface by invoking an operation on the proxy for the object: to read the current time, the client invokes the `getTime` operation; to set the current time, the client invokes the `setTime` operation, passing an argument of type `TimeOfDay`.

Invoking an operation on a proxy instructs the Ice run time to send a message to the target object. The target object can be in another address space or can be collocated (in the same process) as the caller — the location of the target object is transparent to the client. If the target object is in another (possibly remote) address space, the Ice run time invokes the operation via a remote procedure call; if the target is collocated with the client, the Ice run time bypasses the network stack altogether to deliver the request more efficiently.

You can think of an interface definition as the equivalent of the public part of a C++ class definition or as the equivalent of a Java interface, and of operation definitions as (virtual) member functions. Note that nothing but operation definitions are allowed to appear inside an interface definition. In particular, you cannot define a type, an exception, or a data member inside an interface. This does not mean that your object implementation cannot contain state — it can, but how that state is implemented (in the form of data members or otherwise) is hidden from the client and, therefore, need not appear in the object's interface definition.

An Ice object has exactly one (most derived) Slice interface type (or class type). Of course, you can create multiple Ice objects that have the same type; to draw the analogy with C++, a Slice interface corresponds to a C++ class definition, whereas an Ice object corresponds to a C++ class instance (but Ice objects can be implemented in multiple different address spaces).

Ice also provides multiple interfaces via a feature called *facets*.

A Slice interface defines the smallest grain of distribution in Ice: each Ice object has a unique identity (encapsulated in its proxy) that distinguishes it from all other Ice objects; for communication to take place, you must invoke operations on an object's proxy. There is no other notion of an addressable entity in Ice. You cannot, for example, instantiate a Slice structure and have clients manipulate that structure remotely. To make the structure accessible, you must create an interface that allows clients to access the structure.

The partition of an application into interfaces therefore has profound influence on the overall architecture. Distribution boundaries must follow interface (or class) boundaries; you can spread the implementation of interfaces over multiple address spaces (and you can implement multiple interfaces in the same address space), but you cannot implement parts of interfaces in different address spaces.

Topics

- Operations
- User Exceptions
- Run-Time Exceptions
- Proxies for Ice Objects
- Interface Inheritance

See Also

- Classes
- Versioning

Copyright 2016, ZeroC, Inc.

Operations

On this page:

- Parameters and Return Values
- Optional Parameters and Return Values
- Style of Operation Definition
- Overloading Operations
- Idempotent Operations

Parameters and Return Values

An operation definition must contain a return type and zero or more parameter definitions. For example, in the `Clock` interface, the `getTime` operation has a return type of `TimeOfDay` and the `setTime` operation has a return type of `void`. You must use `void` to indicate that an operation returns no value — there is no default return type for Slice operations.

An operation can have one or more input parameters. For example, `setTime` accepts a single input parameter of type `TimeOfDay` called `time`. Of course, you can use multiple input parameters:

Slice

```
interface CircadianRhythm {
    void setSleepPeriod(TimeOfDay startTime, TimeOfDay stopTime);
    // ...
};
```

Note that the parameter name (as for Java) is mandatory. You cannot omit the parameter name, so the following is in error:

Slice

```
interface CircadianRhythm {
    void setSleepPeriod(TimeOfDay, TimeOfDay);  // Error!
    // ...
};
```

By default, parameters are sent from the client to the server, that is, they are input parameters. To pass a value from the server to the client, you can use an output parameter, indicated by the `out` keyword. For example, an alternative way to define the `getTime` operation in the `Clock` interface would be:

Slice

```
void getTime(out TimeOfDay time);
```

This achieves the same thing but uses an output parameter instead of the return value. As with input parameters, you can use multiple output parameters:

Copyright 2016, ZeroC, Inc.

Slice

```
interface CircadianRhythm {
    void setSleepPeriod(TimeOfDay startTime, TimeOfDay stopTime);
    void getSleepPeriod(out TimeOfDay startTime, out TimeOfDay stopTime);
    // ...
};
```

If you have both input and output parameters for an operation, the output parameters must follow the input parameters:

Slice

```
void changeSleepPeriod(    TimeOfDay startTime,         TimeOfDay stopTime,     //
OK
                       out TimeOfDay prevStartTime, out TimeOfDay prevStopTime);

void changeSleepPeriod(out TimeOfDay prevStartTime, out TimeOfDay prevStopTime, //
Error
                           TimeOfDay startTime,         TimeOfDay stopTime);
```

Slice does not support parameters that are both input and output parameters (call by reference). The reason is that, for remote calls, reference parameters do not result in the same savings that one can obtain for call by reference in programming languages. (Data still needs to be copied in both directions and any gains in marshaling efficiency are negligible.) Also, reference (or input-output) parameters result in more complex language mappings, with concomitant increases in code size.

Optional Parameters and Return Values

As of Ice 3.5, an operation's return value and parameters may be declared as optional to indicate that a program can leave their values unset. Parameters not declared as optional are known as *required* parameters; a program must supply legal values for all required parameters. In the discussion below, we use *parameter* to refer to input parameters, output parameters, and return values.

A unique, non-negative integer *tag* must be assigned to each optional parameter:

Slice

```
optional(3) bool example(optional(2) string name, out optional(1) int value);
```

The scope of a tag is limited to its operation and has no effect on other operations.

An operation's signature can include any combination of required and optional parameters, but output parameters still must follow input parameters:

Copyright 2016, ZeroC, Inc.

Slice

```
bool example(string name, optional(3) string referrer, out optional(1) string
promo, out int id);
```

Language mappings specify an API for passing optional parameters and testing whether a parameter is present. Refer to the language mapping sections for more details on the optional parameter API.

A well-behaved program must test for the presence of an optional parameter and not assume that it is always set. Dereferencing an unset optional parameter causes a run-time error.

Style of Operation Definition

As you would expect, language mappings follow the style of operation definition you use in Slice: Slice return types map to programming language return types, and Slice parameters map to programming language parameters.

For operations that return only a single value, it is common to return the value from the operation instead of using an out-parameter. This style maps naturally into all programming languages. Note that, if you use an out-parameter instead, you impose a different API style on the client: most programming languages permit the return value of a function to be ignored whereas it is typically not possible to ignore an output parameter.

For operations that return multiple values, it is common to return all values as out-parameters and to use a return type of `void`. However, the rule is not all that clear-cut because operations with multiple output values can have one particular value that is considered more "important" than the remainder. A common example of this is an iterator operation that returns items from a collection one-by-one:

Slice

```
bool next(out RecordType r);
```

The `next` operation returns two values: the record that was retrieved and a Boolean to indicate the end-of-collection condition. (If the return value is `false`, the end of the collection has been reached and the parameter `r` has an undefined value.) This style of definition can be useful because it naturally fits into the way programmers write control structures. For example:

```
while (next(record))
    // Process record...

if (next(record))
    // Got a valid record...
```

Overloading Operations

Slice does not support any form of overloading of operations. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
interface CircadianRhythm {
    void modify(TimeOfDay startTime, TimeOfDay endTime);
    void modify(    TimeOfDay startTime,        // Error
                    TimeOfDay endTime,
                out timeOfDay prevStartTime,
                out TimeOfDay prevEndTime);
};
```

Operations in the same interface must have different names, regardless of what type and number of parameters they have. This restriction exists because overloaded functions cannot sensibly be mapped to languages without built-in support for overloading.

Name mangling is not an option in this case: while it works fine for compilers, it is unacceptable to humans.

Idempotent Operations

Some operations, such as `getTime` in the `Clock` interface, do not modify the state of the object they operate on. They are the conceptual equivalent of C++ `const` member functions. Similary, `setTime` does modify the state of the object, but is idempotent. You can indicate this in Slice as follows:

Slice

```
interface Clock {
    idempotent TimeOfDay getTime();
    idempotent void setTime(TimeOfDay time);
};
```

This marks the `getTime` and `setTime` operations as idempotent. An operation is idempotent if two successive invocations of the operation have the same effect as a single invocation. For example, `x = 1;` is an idempotent operation because it does not matter whether it is executed once or twice — either way, `x` ends up with the value 1. On the other hand, `x += 1;` is not an idempotent operation because executing it twice results in a different value for `x` than executing it once. Obviously, any read-only operation is idempotent.

The `idempotent` keyword is useful because it allows the Ice run time to be more aggressive when performing automatic retries to recover from errors. Specifically, Ice guarantees *at-most-once* semantics for operation invocations:

- For normal (not idempotent) operations, the Ice run time has to be conservative about how it deals with errors. For example, if a client sends an operation invocation to a server and then loses connectivity, there is no way for the client-side run time to find out whether the request it sent actually made it to the server. This means that the run time cannot attempt to recover from the error by re-establishing a connection and sending the request a second time because that could cause the operation to be invoked a second time and violate at-most-once semantics; the run time has no option but to report the error to the application.
- For `idempotent` operations, on the other hand, the client-side run time can attempt to re-establish a connection to the server and safely send the failed request a second time. If the server can be reached on the second attempt, everything is fine and the application never notices the (temporary) failure. Only if the second attempt fails need the run time report the error back to the application. (The number of retries can be increased with an Ice configuration parameter.)

See Also

- Interfaces, Operations, and Exceptions
- User Exceptions
- Run-Time Exceptions
- Proxies for Ice Objects
- Interface Inheritance

Copyright 2016, ZeroC, Inc.

- Automatic Retries
- Optional Values

Copyright 2016, ZeroC, Inc.

User Exceptions

On this page:

- User Exception Syntax and Semantics
  - Default Values for User Exception Members
  - Declaring User Exceptions in Operations
  - Restrictions for User Exceptions
- User Exception Inheritance

User Exception Syntax and Semantics

Looking at the `setTime` operation in the `Clock` interface, we find a potential problem: given that the `TimeOfDay` structure uses `short` as the type of each field, what will happen if a client invokes the `setTime` operation and passes a `TimeOfDay` value with meaningless field values, such as `-199` for the minute field, or `42` for the hour? Obviously, it would be nice to provide some indication to the caller that this is meaningless. Slice allows you to define user exceptions to indicate error conditions to the client. For example:

Slice

```
exception Error {}; // Empty exceptions are legal

exception RangeError {
    TimeOfDay errorTime;
    TimeOfDay minTime;
    TimeOfDay maxTime;
};
```

A user exception is much like a structure in that it contains a number of data members. However, unlike structures, exceptions can have zero data members, that is, be empty. Like classes, user exceptions support inheritance and may include optional data members.

Default Values for User Exception Members

You can specify a default value for an exception data member that has one of the following types:

- An integral type (`byte`, `short`, `int`, `long`)
- A floating point type (`float` or `double`)
- `string`
- `bool`
- `enum`

For example:

Slice

```
exception RangeError {
    TimeOfDay errorTime;
    TimeOfDay minTime;
    TimeOfDay maxTime;
    string reason = "out of range";
};
```

The legal syntax for literal values is the same as for Slice constants, and you may also use a constant as a default value. The language mapping guarantees that data members are initialized to their declared default values using a language-specific mechanism.

Declaring User Exceptions in Operations

Copyright 2016, ZeroC, Inc.

Exceptions allow you to return an arbitrary amount of error information to the client if an error condition arises in the implementation of an operation. Operations use an exception specification to indicate the exceptions that may be returned to the client:

Slice

```
interface Clock {
    idempotent TimeOfDay getTime();
    idempotent void setTime(TimeOfDay time)
        throws RangeError, Error;
};
```

This definition indicates that the `setTime` operation may throw either a `RangeError` or an `Error` user exception (and no other type of exception). If the client receives a `RangeError` exception, the exception contains the `TimeOfDay` value that was passed to `setTime` and caused the error (in the `errorTime` member), as well as the minimum and maximum time values that can be used (in the `minTime` and `maxTime` members). If `setTime` failed because of an error not caused by an illegal parameter value, it throws `Error`. Obviously, because `Error` does not have data members, the client will have no idea what exactly it was that went wrong — it simply knows that the operation did not work.

An operation can throw only those user exceptions that are listed in its exception specification. If, at run time, the implementation of an operation throws an exception that is not listed in its exception specification, the client receives a run-time exception) to indicate that the operation did something illegal. To indicate that an operation does not throw any user exception, simply omit the exception specification. (There is no empty exception specification in Slice.)

Restrictions for User Exceptions

Exceptions are not first-class data types and first-class data types are not exceptions:

- You cannot pass an exception as a parameter value.
- You cannot use an exception as the type of a data member.
- You cannot use an exception as the element type of a sequence.
- You cannot use an exception as the key or value type of a dictionary.
- You cannot throw a value of non-exception type (such as a value of type `int` or `string`).

The reason for these restrictions is that some implementation languages use a specific and separate type for exceptions (in the same way as Slice does). For such languages, it would be difficult to map exceptions if they could be used as an ordinary data type. (C++ is somewhat unusual among programming languages by allowing arbitrary types to be used as exceptions.)

User Exception Inheritance

Exceptions support inheritance. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
exception ErrorBase {
    string reason;
};

enum RTError {
    DivideByZero, NegativeRoot, IllegalNull /* ... */
};

exception RuntimeError extends ErrorBase {
    RTError err;
};

enum LError { ValueOutOfRange, ValuesInconsistent, /* ... */ };

exception LogicError extends ErrorBase {
    LError err;
};

exception RangeError extends LogicError {
    TimeOfDay errorTime;
    TimeOfDay minTime;
    TimeOfDay maxTime;
};
```

These definitions set up a simple exception hierarchy:

- `ErrorBase` is at the root of the tree and contains a string explaining the cause of the error.
- Derived from `ErrorBase` are `RuntimeError` and `LogicError`. Each of these exceptions contains an enumerated value that further categorizes the error.
- Finally, `RangeError` is derived from `LogicError` and reports the details of the specific error.

Setting up exception hierarchies such as this not only helps to create a more readable specification because errors are categorized, but also can be used at the language level to good advantage. For example, the Slice C++ mapping preserves the exception hierarchy so you can catch exceptions generically as a base exception, or set up exception handlers to deal with specific exceptions.

Looking at the exception hierarchy, it is not clear whether, at run time, the application will only throw most derived exceptions, such as `RangeError`, or if it will also throw base exceptions, such as `LogicError`, `RuntimeError`, and `ErrorBase`. If you want to indicate that a base exception, interface, or class is abstract (will not be instantiated), you can add a comment to that effect.

Note that, if the exception specification of an operation indicates a specific exception type, at run time, the implementation of the operation may also throw more derived exceptions. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
exception Base {
    // ...
};

exception Derived extends Base {
    // ...
};

interface Example {
    void op() throws Base;      // May throw Base or Derived
};
```

In this example, `op` may throw a `Base` or a `Derived` exception, that is, any exception that is compatible with the exception types listed in the exception specification can be thrown at run time.

As a system evolves, it is quite common for new, derived exceptions to be added to an existing hierarchy. Assume that we initially construct clients and server with the following definitions:

Slice

```
exception Error {
    // ...
};

interface Application {
    void doSomething() throws Error;
};
```

Also assume that a large number of clients are deployed in field, that is, when you upgrade the system, you cannot easily upgrade all the clients. As the application evolves, a new exception is added to the system and the server is redeployed with the new definition:

Slice

```
exception Error {
    // ...
};

exception FatalApplicationError extends Error {
    // ...
};

interface Application {
    void doSomething() throws Error;
};
```


Copyright 2016, ZeroC, Inc.

This raises the question of what should happen if the server throws a `FatalApplicationError` from `doSomething`. The answer depends whether the client was built using the old or the updated definition:

- If the client was built using the same definition as the server, it simply receives a `FatalApplicationError`.
- If the client was built with the original definition, that client has no knowledge that `FatalApplicationError` even exists. In this case, the Ice run time automatically slices the exception to the most-derived type that is understood by the receiver (`Error`, in this case) and discards the information that is specific to the derived part of the exception. (This is exactly analogous to catching C++ exceptions by value — the exception is sliced to the type used in the `catch`-clause.)

Exceptions support single inheritance only. (Multiple inheritance would be difficult to map into many programming languages.)

See Also

- Constants and Literals
- Operations
- Run-Time Exceptions
- Proxies for Ice Objects
- Interface Inheritance
- Optional Data Members

Copyright 2016, ZeroC, Inc.

Run-Time Exceptions

In addition to any user exceptions that are listed in an operation's exception specification, an operation can also throw Ice *run-time exceptions*. Run-time exceptions are predefined exceptions that indicate platform-related run-time errors. For example, if a networking error interrupts communication between client and server, the client is informed of this by a run-time exception, such as `ConnectTimeoutException` or `SocketException`.

The exception specification of an operation must not list any run-time exceptions. (It is understood that all operations can raise run-time exceptions and you are not allowed to restate that.)

On this page:

- Inheritance Hierarchy for Exceptions
- Local Versus Remote Exceptions
  - Common Exceptions
    - ObjectNotExistException
    - FacetNotExistException
    - OperationNotExistException
  - Unknown Exceptions
    - UnknownUserException
    - UnknownLocalException
    - UnknownException

Inheritance Hierarchy for Exceptions

All the Ice run-time and user exceptions are arranged in an inheritance hierarchy, as shown below:

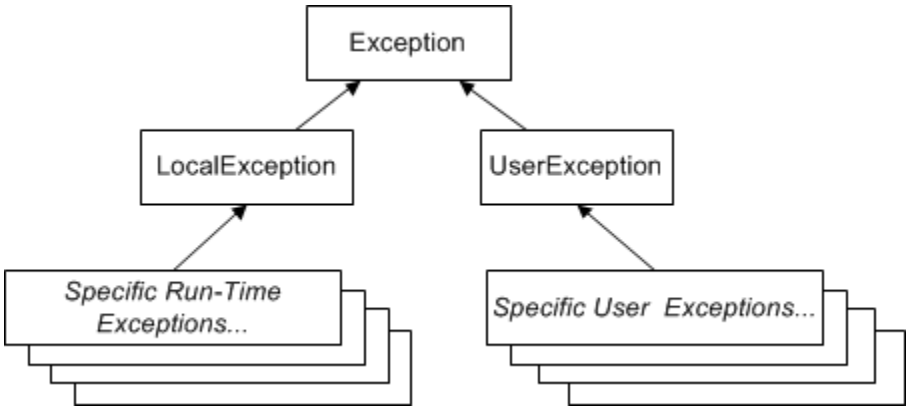

Inheritance structure for exceptions.

`Ice::Exception` is at the root of the inheritance hierarchy. Derived from that are the (abstract) types `Ice::LocalException` and `Ice::UserException`. In turn, all run-time exceptions are derived from `Ice::LocalException`, and all user exceptions are derived from `Ice::UserException`.

This figures shows the complete hierarchy of the Ice run-time exceptions:

Copyright 2016, ZeroC, Inc.

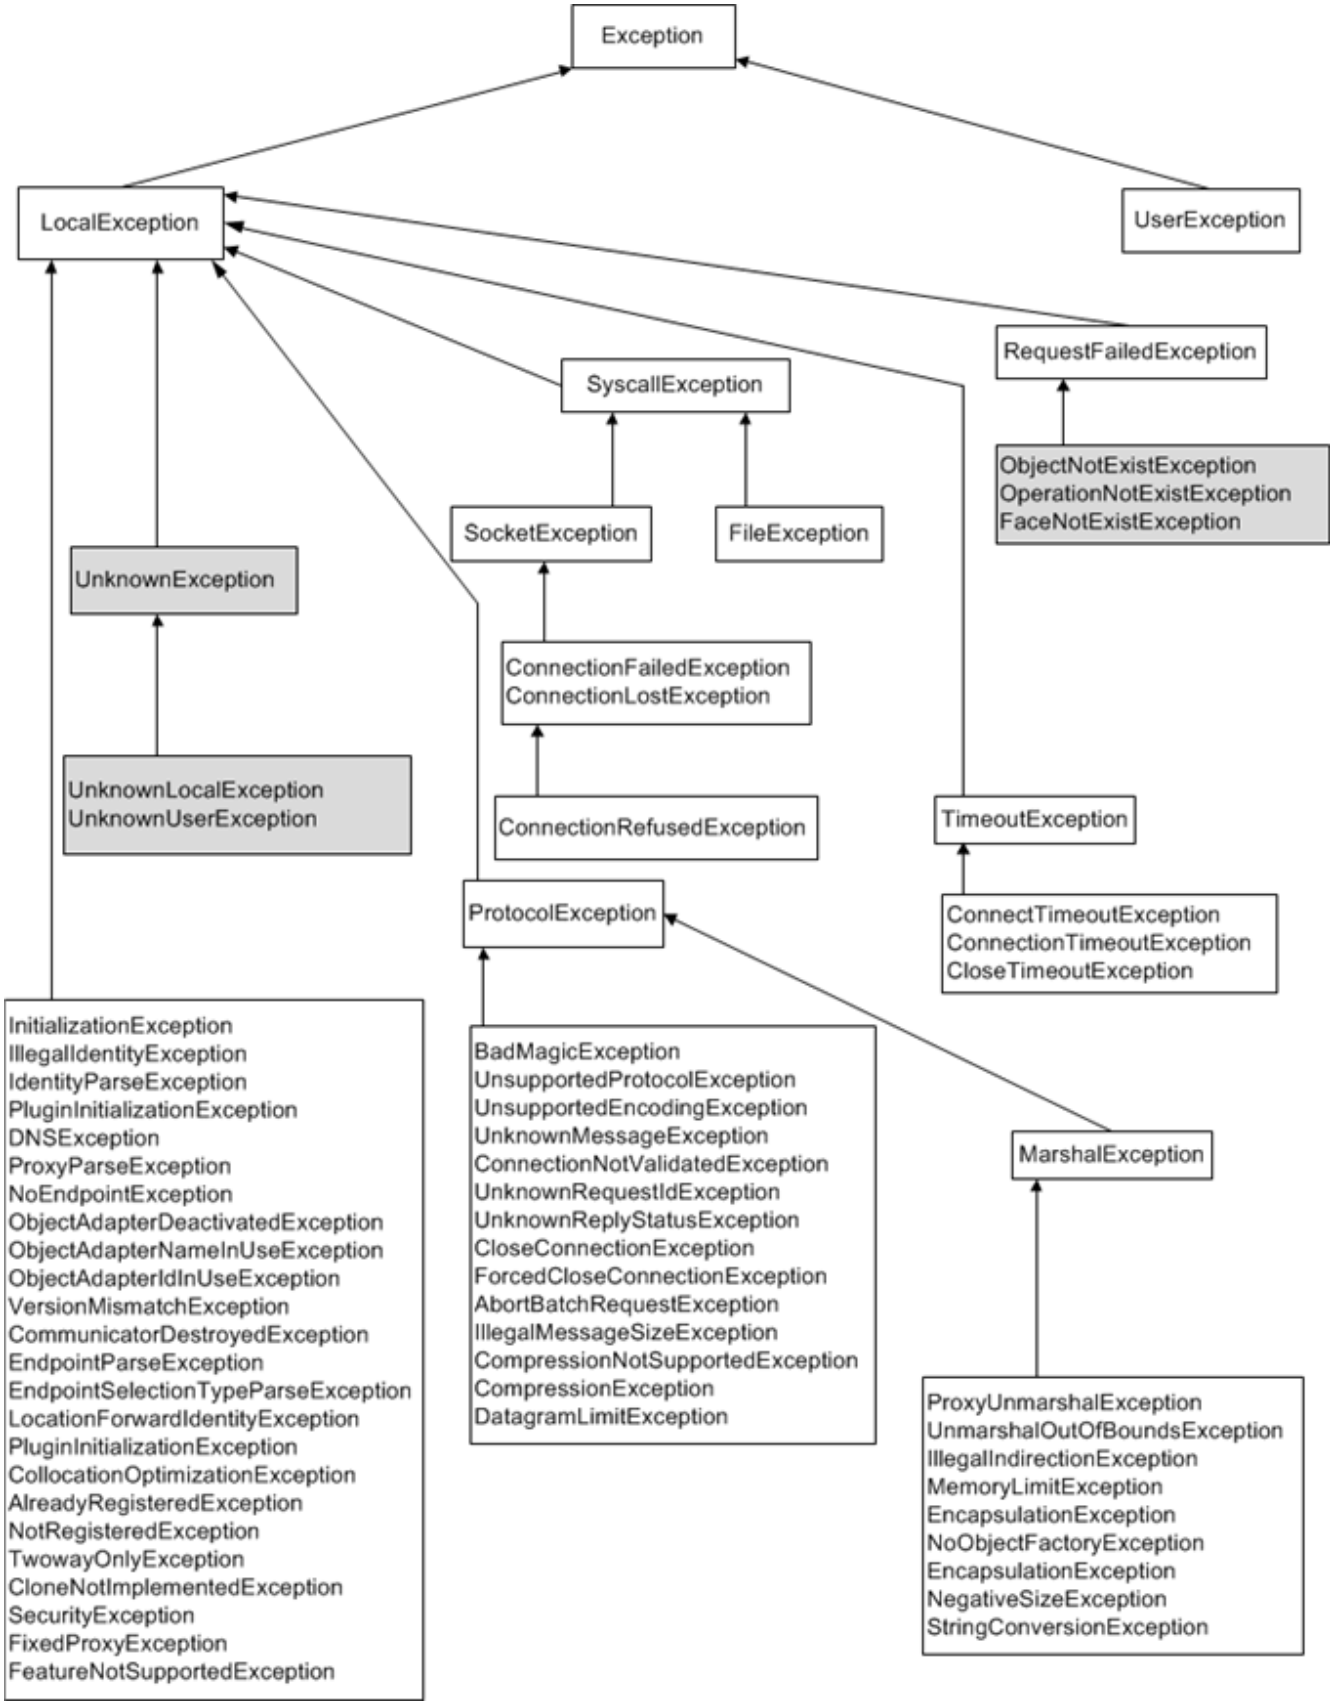

Ice run-time exception hierarchy. (Shaded exceptions can be sent by the server.)

Copyright 2016, ZeroC, Inc.

We use the Unified Modeling Language (UML) for the object model diagrams (see [1] and [2] for details).

Note that Ice run-time exception hierarchy groups several exceptions into a single box to save space (which, strictly, is incorrect UML syntax). Also note that some run-time exceptions have data members, which, for brevity, we have omitted in the Ice run-time exception hierarchy. These data members provide additional information about the precise cause of an error.

Many of the run-time exceptions have self-explanatory names, such as `MemoryLimitException`. Others indicate problems in the Ice run time, such as `EncapsulationException`. Still others can arise only through application programming errors, such as `TwowayOnlyException`. In practice, you will likely never see most of these exceptions. However, there are a few run-time exceptions you will encounter and whose meaning you should know.

Local Versus Remote Exceptions

Common Exceptions

Most error conditions are detected on the client side. For example, if an attempt to contact a server fails, the client-side run time raises a `ConnectTimeoutException`. However, there are three specific error conditions (shown as shaded in the Ice run-time exception hierarchy diagram) that are detected by the server and made known explicitly to the client-side run time via the Ice protocol: `ObjectNotExistException`, `FacetNotExistException`, and `OperationNotExistException`.

`ObjectNotExistException`

This exception indicates that a request was delivered to the server but the server could not locate a servant with the identity that is embedded in the proxy. In other words, the server could not find an object to dispatch the request to.

An `ObjectNotExistException` is a death certificate: it indicates that the target object in the server does not exist.

The Ice run time raises `ObjectNotExistException` only if there are no facets in existence with a matching identity; otherwise, it raises `FacetNotExistException`.

Most likely, this is the case because the object existed some time in the past and has since been destroyed, but the same exception is also raised if a client uses a proxy with the identity of an object that has never been created. If you receive this exception, you are expected to clean up whatever resources you might have allocated that relate to the specific object for which you receive this exception.

`FacetNotExistException`

The client attempted to contact a non-existent facets of an object, that is, the server has at least one servant with the given identity, but no servant with a matching facet name.

`OperationNotExistException`

This exception is raised if the server could locate an object with the correct identity but, on attempting to dispatch the client's operation invocation, the server found that the target object does not have such an operation. You will see this exception in only two cases:

- You have used an unchecked down-cast on a proxy of the incorrect type.
- Client and server have been built with Slice definitions for an interface that disagree with each other, that is, the client was built with an interface definition for the object that indicates that an operation exists, but the server was built with a different version of the interface definition in which the operation is absent.

Unknown Exceptions

Any error condition on the server side that is not described by one of the three preceding exceptions is made known to the client as one of three generic exceptions (shown as shaded in the Ice run-time exception hierarchy figure diagram): `UnknownUserException`, `UnknownLocalException`, or `UnknownException`.

`UnknownUserException`

This exception indicates that an operation implementation has thrown a Slice exception that is not declared in the operation's exception specification (and is not derived from one of the exceptions in the operation's exception specification).

`UnknownLocalException`

Copyright 2016, ZeroC, Inc.

If an operation implementation raises a run-time exception other than `ObjectNotExistException`, `FacetNotExistException`, or `OperationNotExistException` (such as a `NotRegisteredException`), the client receives an `UnknownLocalException`. In other words, the Ice protocol does not transmit the exact exception that was encountered in the server, but simply returns a bit to the client in the reply to indicate that the server encountered a run-time exception.

A common cause for a client receiving an `UnknownLocalException` is failure to catch and handle all exceptions in the server. For example, if the implementation of an operation encounters an exception it does not handle, the exception propagates all the way up the call stack until the stack is unwound to the point where the Ice run time invoked the operation. The Ice run time catches all Ice exceptions that "escape" from an operation invocation and returns them to the client as an `UnknownLocalException`.

`UnknownException`

An operation has thrown a non-Ice exception. For example, if the operation in the server throws a C++ exception, such as a `char*`, or a Java exception, such as a `ClassCastException`, the client receives an `UnknownException`.

All other run-time exceptions (not shaded in the Ice run-time exception hierarchy) are detected by the client-side run time and are raised locally.

It is possible for the implementation of an operation to throw Ice run-time exceptions (as well as user exceptions). For example, if a client holds a proxy to an object that no longer exists in the server, your server application code is required to throw an `ObjectNotExistException`. If you do throw run-time exceptions from your application code, you should take care to throw a run-time exception only if appropriate, that is, do not use run-time exceptions to indicate something that really should be a user exception. Doing so can be very confusing to the client: if the application "hijacks" some run-time exceptions for its own purposes, the client can no longer decide whether the exception was thrown by the Ice run time or by the server application code. This can make debugging very difficult.

See Also

- User Exceptions
- Interfaces, Operations, and Exceptions
- Operations
- Proxies for Ice Objects
- Interface Inheritance
- Versioning

References

1. Booch, G., et al. 1998. *Unified Modeling Language User Guide*. Reading, MA: Addison-Wesley.
2. Object Management Group. 2001. *Unified Modeling Language Specification*. Framingham, MA: Object Management Group.

Copyright 2016, ZeroC, Inc.

Proxies for Ice Objects

Building on the `Clock` example, we can create definitions for a world-time server:

Slice

```
exception GenericError {
    string reason;
};

struct TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
};

exception BadTimeVal extends GenericError {};

interface Clock {
    idempotent TimeOfDay getTime();
    idempotent void setTime(TimeOfDay time) throws BadTimeVal;
};

dictionary<string, Clock*> TimeMap; // Time zone name to clock map

exception BadZoneName extends GenericError {};

interface WorldTime {
    idempotent void addZone(string zoneName, Clock* zoneClock);
    void removeZone(string zoneName) throws BadZoneName;
    idempotent Clock* findZone(string zoneName) throws BadZoneName;
    idempotent TimeMap listZones();
    idempotent void setZones(TimeMap zones);
};
```

The `WorldTime` interface acts as a collection manager for clocks, one for each time zone. In other words, the `WorldTime` interface manages a collection of pairs. The first member of each pair is a time zone name; the second member of the pair is the clock that provides the time for that zone. The interface contains operations that permit you to add or remove a clock from the map (`addZone` and `removeZone`), to search for a particular time zone by name (`findZone`), and to read or write the entire map (`listZones` and `setZones`).

The `WorldTime` example illustrates an important Slice concept: note that `addZone` accepts a parameter of type `Clock*` and `findZone` ret urns a parameter of type `Clock*`. In other words, interfaces are types in their own right and can be passed as parameters. The `*` operator is known as the *proxy operator*. Its left-hand argument must be an interface (or class) and its return type is a proxy. A proxy is like a pointer that can denote an object. The semantics of proxies are very much like those of C++ class instance pointers:

- A proxy can be null.
- A proxy can dangle (point at an object that is no longer there).
- Operations dispatched via a proxy use late binding: if the actual run-time type of the object denoted by the proxy is more derived than the proxy's type, the implementation of the most-derived interface will be invoked.

When a client passes a `Clock` proxy to the `addZone` operation, the proxy denotes an actual `Clock` object in a server. The `Clock` *Ice object* denoted by that proxy may be implemented in the same server process as the `WorldTime` interface, or in a different server process. Where the `Clock` object is physically implemented matters neither to the client nor to the server implementing the `WorldTime` interface; if either invokes an operation on a particular clock, such as `getTime`, an RPC call is sent to whatever server implements that particular clock. In other words, a proxy acts as a local "ambassador" for the remote object; invoking an operation on the proxy forwards the invocation to the actual object implementation. If the object implementation is in a different address space, this results in a remote procedure call; if the object implementation is collocated in the same address space, the Ice run time may optimize the invocation.

Copyright 2016, ZeroC, Inc.

Note that proxies also act very much like pointers in their sharing semantics: if two clients have a proxy to the same object, a state change made by one client (such as setting the time) will be visible to the other client.

Proxies are strongly typed (at least for statically typed languages, such as C++ and Java). This means that you cannot pass something other than a `Clock` proxy to the `addZone` operation; attempts to do so are rejected at compile time.

See Also

- Classes
- Interfaces, Operations, and Exceptions
- User Exceptions
- Run-Time Exceptions
- Interface Inheritance

Copyright 2016, ZeroC, Inc.

Interface Inheritance

On this page:

- Interface Inheritance
- Interface Inheritance Limitations
- Implicit Inheritance from Object
- Null Proxies
- Self-Referential Interfaces
- Empty Interfaces
- Interface Versus Implementation Inheritance

Interface Inheritance

Interfaces support inheritance. For example, we could extend our world-time server to support the concept of an alarm clock:

Slice

```
interface AlarmClock extends Clock {
    idempotent TimeOfDay getAlarmTime();
    idempotent void setAlarmTime(TimeOfDay alarmTime)
        throws BadTimeVal;
};
```

The semantics of this are the same as for C++ or Java: `AlarmClock` is a subtype of `Clock` and an `AlarmClock` proxy can be substituted wherever a `Clock` proxy is expected. Obviously, an `AlarmClock` supports the same `getTime` and `setTime` operations as a `Clock` but also supports the `getAlarmTime` and `setAlarmTime` operations.

Multiple interface inheritance is also possible. For example, we can construct a radio alarm clock as follows:

Slice

```
interface Radio {
    void setFrequency(long hertz) throws GenericError;
    void setVolume(long dB) throws GenericError;
};

enum AlarmMode { RadioAlarm, BeepAlarm };

interface RadioClock extends Radio, AlarmClock {
    void setMode(AlarmMode mode);
    AlarmMode getMode();
};
```

`RadioClock` extends both `Radio` and `AlarmClock` and can therefore be passed where a `Radio`, an `AlarmClock`, or a `Clock` is expected. The inheritance diagram for this definition looks as follows:

Copyright 2016, ZeroC, Inc.

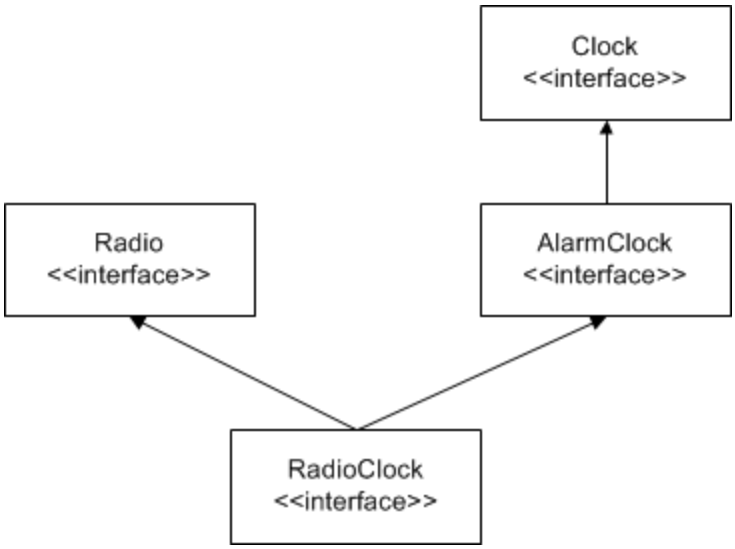

Inheritance diagram for `RadioClock`.

Interfaces that inherit from more than one base interface may share a common base interface. For example, the following definition is legal:

Slice

```
interface B { /* ... */ };
interface I1 extends B { /* ... */ };
interface I2 extends B { /* ... */ };
interface D extends I1, I2 { /* ... */ };
```

This definition results in the familiar diamond shape:

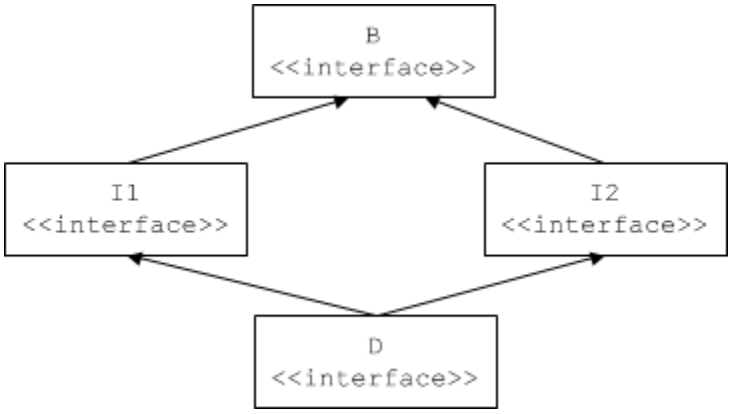

Diamond-shaped inheritance.

Interface Inheritance Limitations

If an interface uses multiple inheritance, it must not inherit the same operation name from more than one base interface. For example, the following definition is illegal:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Clock {
    void set(TimeOfDay time);                   // set time
};

interface Radio {
    void set(long hertz);                       // set frequency
};

interface RadioClock extends Radio, Clock {     // Illegal!
    // ...
};
```

This definition is illegal because `RadioClock` inherits two `set` operations, `Radio::set` and `Clock::set`. The Slice compiler makes this illegal because (unlike C++) many programming languages do not have a built-in facility for disambiguating the different operations. In Slice, the simple rule is that all inherited operations must have unique names. (In practice, this is rarely a problem because inheritance is rarely added to an interface hierarchy "after the fact". To avoid accidental clashes, we suggest that you use descriptive operation names, such as `setTime` and `setFrequency`. This makes accidental name clashes less likely.)

Implicit Inheritance from Object

All Slice interfaces are ultimately derived from `Object`. For example, the inheritance hierarchy would be shown more correctly as:

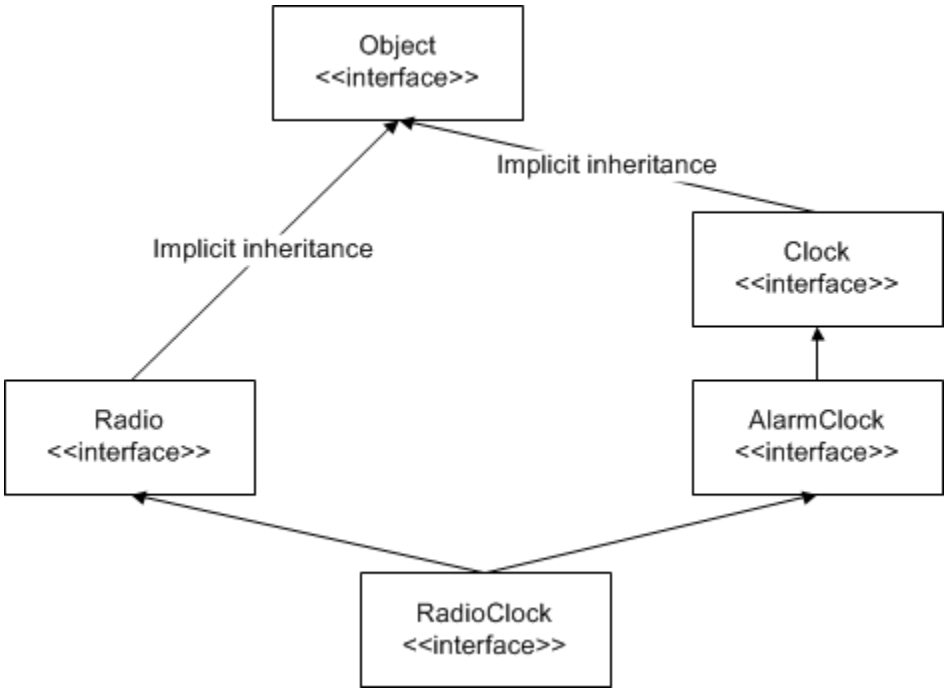

Implicit inheritance from `Object`.

Because all interfaces have a common base interface, we can pass any type of interface as that type. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
interface ProxyStore {
    idempotent void putProxy(string name, Object* o);
    idempotent Object* getProxy(string name);
};
```

`Object` is a Slice keyword (note the capitalization) that denotes the root type of the inheritance hierarchy. The `ProxyStore` interface is a generic proxy storage facility: the client can call `putProxy` to add a proxy of any type under a given name and later retrieve that proxy again by calling `getProxy` and supplying that name. The ability to generically store proxies in this fashion allows us to build general-purpose facilities, such as a naming service that can store proxies and deliver them to clients. Such a service, in turn, allows us to avoid hard-coding proxy details into clients and servers.

Inheritance from type `Object` is always implicit. For example, the following Slice definition is illegal:

Slice

```
interface MyInterface extends Object { /* ... */ }; // Error!
```

It is understood that all interfaces inherit from type `Object`; you are not allowed to restate that.

Type `Object` is mapped to an abstract type by the various language mappings, so you cannot instantiate an Ice object of that type.

Null Proxies

Looking at the `ProxyStore` interface once more, we notice that `getProxy` does not have an exception specification. The question then is what should happen if a client calls `getProxy` with a name under which no proxy is stored? Obviously, we could add an exception to indicate this condition to `getProxy`. However, another option is to return a *null proxy*. Ice has the built-in notion of a null proxy, which is a proxy that "points nowhere". When such a proxy is returned to the client, the client can test the value of the returned proxy to check whether it is null or denotes a valid object.

A more interesting question is: "which approach is more appropriate, throwing an exception or returning a null proxy?" The answer depends on the expected usage pattern of an interface. For example, if, in normal operation, you do not expect clients to call `getProxy` with a non-existent name, it is better to throw an exception. (This is probably the case for our `ProxyStore` interface: the fact that there is no `list` operation makes it clear that clients are expected to know which names are in use.)

On the other hand, if you expect that clients will occasionally try to look up something that is not there, it is better to return a null proxy. The reason is that throwing an exception breaks the normal flow of control in the client and requires special handling code. This means that you should throw exceptions only in exceptional circumstances. For example, throwing an exception if a database lookup returns an empty result set is wrong; it is expected and normal that a result set is occasionally empty.

It is worth paying attention to such design issues: well-designed interfaces that get these details right are easier to use and easier to understand. Not only do such interfaces make life easier for client developers, they also make it less likely that latent bugs cause problems later.

Self-Referential Interfaces

Proxies have pointer semantics, so we can define self-referential interfaces. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Link {
    idempotent SomeType getValue();
    idempotent Link* next();
};
```

The `Link` interface contains a `next` operation that returns a proxy to a `Link` interface. Obviously, this can be used to create a chain of interfaces; the final link in the chain returns a null proxy from its `next` operation.

Empty Interfaces

The following Slice definition is legal:

Slice

```
interface Empty {};
```

The Slice compiler will compile this definition without complaint. An interesting question is: "why would I need an empty interface?" In most cases, empty interfaces are an indication of design errors. Here is one example:

Slice

```
interface ThingBase {};

interface Thing1 extends ThingBase {
    // Operations here...
};

interface Thing2 extends ThingBase {
    // Operations here...
};
```

Looking at this definition, we can make two observations:

- `Thing1` and `Thing2` have a common base and are therefore related.
- Whatever is common to `Thing1` and `Thing2` can be found in interface `ThingBase`.

Of course, looking at `ThingBase`, we find that `Thing1` and `Thing2` do not share any operations at all because `ThingBase` is empty. Given that we are using an object-oriented paradigm, this is definitely strange: in the object-oriented model, the *only* way to communicate with an object is to send a message to the object. But, to send a message, we need an operation. Given that `ThingBase` has no operations, we cannot send a message to it, and it follows that `Thing1` and `Thing2` are *not* related because they have no common operations. But of course, seeing that `Thing1` and `Thing2` have a common base, we conclude that they *are* related, otherwise the common base would not exist. At this point, most programmers begin to scratch their head and wonder what is going on here.

One common use of the above design is a desire to treat `Thing1` and `Thing2` polymorphically. For example, we might continue the previous definition as follows:

Copyright 2016, ZeroC, Inc.

Slice

```
interface ThingUser {
    void putThing(ThingBase* thing);
};
```

Now the purpose of having the common base becomes clear: we want to be able to pass both `Thing1` and `Thing2` proxies to `putThing`. Does this justify the empty base interface? To answer this question, we need to think about what happens in the implementation of `putThing`. Obviously, `putThing` cannot possibly invoke an operation on a `ThingBase` because there are no operations. This means that `putThing` can do one of two things:

1. putThing can simply remember the value of `thing`.
2. putThing can try to down-cast to either `Thing1` or `Thing2` and then invoke an operation. The pseudo-code for the implementation of `putThing` would look something like this:

```
void putThing(ThingBase thing)
{
    if (is_a(Thing1, thing)) {
        // Do something with Thing1...
    } else if (is_a(Thing2, thing)) {
        // Do something with Thing2...
    } else {
        // Might be a ThingBase?
        // ...
    }
}
```

   The implementation tries to down-cast its argument to each possible type in turn until it has found the actual run-time type of the argument. Of course, any object-oriented text book worth its price will tell you that this is an abuse of inheritance and leads to maintenance problems.

If you find yourself writing operations such as `putThing` that rely on artificial base interfaces, ask yourself whether you really need to do things this way. For example, a more appropriate design might be:

Slice

```
interface Thing1 {
    // Operations here...
};

interface Thing2 {
    // Operations here...
};

interface ThingUser {
    void putThing1(Thing1* thing);
    void putThing2(Thing2* thing);
};
```


Copyright 2016, ZeroC, Inc.

With this design, `Thing1` and `Thing2` are not related, and `ThingUser` offers a separate operation for each type of proxy. The implementation of these operations does not need to use any down-casts, and all is well in our object-oriented world.

Another common use of empty base interfaces is the following:

Slice

```
interface PersistentObject {};

interface Thing1 extends PersistentObject {
    // Operations here...
};

interface Thing2 extends PersistentObject {
    // Operations here...
};
```

Clearly, the intent of this design is to place persistence functionality into the `PersistentObject` base *implementation* and require objects that want to have persistent state to inherit from `PersistentObject`. On the face of things, this is reasonable: after all, using inheritance in this way is a well-established design pattern, so what can possibly be wrong with it? As it turns out, there are a number of things that are wrong with this design:

- The above inheritance hierarchy is used to add *behavior* to `Thing1` and `Thing2`. However, in a strict OO model, behavior can be invoked only by sending messages. But, because `PersistentObject` has no operations, no messages can be sent.
  This raises the question of how the implementation of `PersistentObject` actually goes about doing its job; presumably, it knows something about the implementation (that is, the internal state) of `Thing1` and `Thing2`, so it can write that state into a database. But, if so, `PersistentObject`, `Thing1`, and `Thing2` can no longer be implemented in different address spaces because, in that case, `PersistentObject` can no longer get at the state of `Thing1` and `Thing2`.
  Alternatively, `Thing1` and `Thing2` use some functionality provided by `PersistentObject` in order to make their internal state persistent. But `PersistentObject` does not have any operations, so how would `Thing1` and `Thing2` actually go about achieving this? Again, the only way that can work is if `PersistentObject`, `Thing1`, and `Thing2` are implemented in a single address space and share implementation state behind the scenes, meaning that they cannot be implemented in different address spaces.
- The above inheritance hierarchy splits the world into two halves, one containing persistent objects and one containing non-persistent ones. This has far-reaching ramifications:
  - Suppose you have an existing application with already implemented, non-persistent objects. Requirements change over time and you find that you now would like to make some of your objects persistent. With the above design, you cannot do this unless you change the type of your objects because they now must inherit from `PersistentObject`. Of course, this is extremely bad news: not only do you have to change the implementation of your objects in the server, you also need to locate and update all the clients that are currently using your objects because they suddenly have a completely new type. What is worse, there is no way to keep things backward compatible: either all clients change with the server, or none of them do. It is impossible for some clients to remain "unupgraded".
  - The design does not scale to multiple features. Imagine that we have a number of additional behaviors that objects can inherit, such as serialization, fault-tolerance, persistence, and the ability to be searched by a search engine. We quickly end up in a mess of multiple inheritance. What is worse, each possible combination of features creates a completely separate type hierarchy. This means that you can no longer write operations that generically operate on a number of object types. For example, you cannot pass a persistent object to something that expects a non-persistent object, *even if the receiver of the object does not care about the persistence aspects of the object*. This quickly leads to fragmented and hard-to-maintain type systems. Before long, you will either find yourself rewriting your application or end up with something that is both difficult to use and difficult to maintain.

The foregoing discussion will hopefully serve as a warning: Slice is an *interface* definition language that has nothing to do with *implementation* (but empty interfaces almost always indicate that implementation state is shared via mechanisms other than defined interfaces). If you find yourself writing an empty interface definition, at least step back and think about the problem at hand; there may be a more appropriate design that expresses your intent more cleanly. If you do decide to go ahead with an empty interface regardless, be aware that, almost certainly, you will lose the ability to later change the distribution of the object model over physical server processes because you cannot place an address space boundary between interfaces that share hidden state.

Interface Versus Implementation Inheritance

Keep in mind that Slice interface inheritance applies only to *interfaces*. In particular, if two interfaces are in an inheritance relationship, this in

Copyright 2016, ZeroC, Inc.

no way implies that the implementations of those interfaces must also inherit from each other. You can choose to use implementation inheritance when you implement your interfaces, but you can also make the implementations independent of each other. (To C++ programmers, this often comes as a surprise because C++ uses implementation inheritance by default, and interface inheritance requires extra effort to implement.)

In summary, Slice inheritance simply establishes type compatibility. It says nothing about how interfaces are implemented and, therefore, keeps implementation choices open to whatever is most appropriate for your application.

See Also

- Interfaces, Operations, and Exceptions
- Operations
- User Exceptions
- Run-Time Exceptions
- Proxies for Ice Objects
- IceGrid

Copyright 2016, ZeroC, Inc.

Classes

In addition to interfaces, Slice permits the definition of classes. Classes are like interfaces in that they can have operations and are like structures in that they can have data members. This leads to hybrid objects that can be treated as interfaces and passed by reference, or can be treated as values and passed by value.

Classes support inheritance and are therefore polymorphic: at run time, you can pass a class instance to an operation as long as the actual class type is derived from the formal parameter type in the operation's signature. This also permits classes to be used as type-safe unions, similarly to Pascal's discriminated variant records.

Topics

- Simple Classes
- Class Inheritance
- Class Inheritance Semantics
- Classes as Unions
- Self-Referential Classes
- Classes Versus Structures
- Classes with Operations
- Architectural Implications of Classes
- Classes Implementing Interfaces
- Class Inheritance Limitations
- Pass-by-Value Versus Pass-by-Reference
- Passing Interfaces by Value
- Classes with Compact Type IDs

Copyright 2016, ZeroC, Inc.

Simple Classes

A Slice class definition is similar to a structure definition, but uses the `class` keyword. For example:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
};
```

Apart from the keyword `class`, this definition is identical to the structure example. You can use a Slice class wherever you can use a Slice structure (but, as we will see shortly, for performance reasons, you should not use a class where a structure is sufficient). Unlike structures, classes can be empty:

Slice

```
class EmptyClass {};    // OK
struct EmptyStruct {};  // Error
```

Much the same design considerations as for empty interfaces apply to empty classes: you should at least stop and rethink your approach before committing yourself to an empty class.

A class can define any number of data members, including optional data members. You can also specify a default value for a data member if its type is one of the following:

- An integral type (`byte`, `short`, `int`, `long`)
- A floating point type (`float` or `double`)
- `string`
- `bool`
- `enum`

For example:

Slice

```
class Location {
    string name;
    Point pt;
    bool display = true;
    string source = "GPS";
};
```

The legal syntax for literal values is the same as for Slice constants, and you may also use a constant as a default value. The language mapping guarantees that data members are initialized to their declared default values using a language-specific mechanism.

See Also

- Structures
- Constants and Literals

Copyright 2016, ZeroC, Inc.

- Optional Data Members

Copyright 2016, ZeroC, Inc.

Class Inheritance

Unlike structures, classes support inheritance. For example:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
};

class DateTime extends TimeOfDay {
    short day;          // 1 - 31
    short month;        // 1 - 12
    short year;         // 1753 onwards
};
```

This example illustrates one major reason for using a class: a class can be extended by inheritance, whereas a structure is not extensible. The previous example defines `DateTime` to extend the `TimeOfDay` class with a date.

If you are puzzled by the comment about the year 1753, search the Web for "1752 date change". The intricacies of calendars for various countries prior to that year can keep you occupied for months...

Classes only support single inheritance. The following is illegal:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
};

class Date {
    short day;
    short month;
    short year;
};

class DateTime extends TimeOfDay, Date {   // Error!
    // ...
};
```

A derived class also cannot redefine a data member of its base class:

Copyright 2016, ZeroC, Inc.

Slice

```
class Base {
    int integer;
};

class Derived extends Base {
    int integer;                // Error, integer redefined
};
```

See Also

- Structures

Copyright 2016, ZeroC, Inc.

Class Inheritance Semantics

Classes use the same pass-by-value semantics as structures. If you pass a class instance to an operation, the class and all its members are passed. The usual type compatibility rules apply: you can pass a derived instance where a base instance is expected. If the receiver has static type knowledge of the actual derived run-time type, it receives the derived instance; otherwise, if the receiver does not have static type knowledge of the derived type, depending on the format used to encode the class, it will either fail to read the instance or slice the instance to the base type.

For an example, suppose we have the following definitions:

Slice

```
// In file Clock.ice:

class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
};

interface Clock {
    TimeOfDay getTime();
    void setTime(TimeOfDay time);
};

// In file DateTime.ice:

#include <Clock.ice>

class DateTime extends TimeOfDay {
    short day;          // 1 - 31
    short month;        // 1 - 12
    short year;         // 1753 onwards
};
```

Because `DateTime` is a sub-class of `TimeOfDay`, the server can return a `DateTime` instance from `getTime`, and the client can pass a `DateTime` instance to `setTime`. In this case, if both client and server are linked to include the code generated for both `Clock.ice` and `DateTime.ice`, they each receive the actual derived `DateTime` instance, that is, the actual run-time type of the instance is preserved.

Contrast this with the case where the server is linked to include the code generated for both `Clock.ice` and `DateTime.ice`, but the client is linked only with the code generated for `Clock.ice`. In other words, the server understands the type `DateTime` and can return a `DateTime` instance from `getTime`, but the client only understands `TimeOfDay`. In this case, there are two possible outcomes depending on the format used by the server to encode the instance:

- with the sliced format, the derived `DateTime` instance returned by the server is sliced to its `TimeOfDay` base type in the client
- with the compact format, `getTime` fails with the `Ice::NoObjectFactoryException` exception

See Design Considerations for Objects for additional information on the sliced and compact formats.

Class hierarchies are useful if you need polymorphic *values* (instead of polymorphic *interfaces*). For example:

Copyright 2016, ZeroC, Inc.

Slice

```
class Shape {
    // Definitions for shapes, such as size, center, etc.
};

class Circle extends Shape {
    // Definitions for circles, such as radius...
};

class Rectangle extends Shape {
    // Definitions for rectangles, such as width and length...
};

sequence<Shape> ShapeSeq;

interface ShapeProcessor {
    void processShapes(ShapeSeq ss);
};
```

Note the definition of `ShapeSeq` and its use as a parameter to the `processShapes` operation: the class hierarchy allows us to pass a polymorphic sequence of shapes (instead of having to define a separate operation for each type of shape).

The receiver of a `ShapeSeq` can iterate over the elements of the sequence and down-cast each element to its actual run-time type. (The receiver can also ask each element for its type ID to determine its type.)

See Also

- Structures
- Type IDs

Copyright 2016, ZeroC, Inc.

Classes as Unions

Slice does not offer a dedicated union construct because it is redundant. By deriving classes from a common base class, you can create the same effect as with a union:

Slice

```
interface ShapeShifter {
    Shape translate(Shape s, long xDistance, long yDistance);
};
```

The parameter `s` of the `translate` operation can be viewed as a union of two members: a `Circle` and a `Rectangle`. The receiver of a `Shape` instance can use the type ID of the instance to decide whether it received a `Circle` or a `Rectangle`. Alternatively, if you want something more along the lines of a conventional discriminated union, you can use the following approach:

Slice

```
class UnionDiscriminator {
    int d;
};

class Member1 extends UnionDiscriminator {
    // d == 1
    string s;
    float f;
};

class Member2 extends UnionDiscriminator {
    // d == 2
    byte b;
    int i;
};
```

With this approach, the `UnionDiscriminator` class provides a discriminator value. The "members" of the union are the classes that are derived from `UnionDiscriminator`. For each derived class, the discriminator takes on a distinct value. The receiver of such a union uses the discriminator value in a `switch` statement to select the active union member.

See Also

- Type IDs

Copyright 2016, ZeroC, Inc.

Self-Referential Classes

Classes can be self-referential.

For example:

Slice

```
class Link {
    SomeType value;
    Link next;
};
```

This looks very similar to the self-referential interface example, but the semantics are very different. Note that `value` and `next` are data members, not operations, and that the type of `next` is `Link` (*not* `Link*`). As you would expect, this forms the same linked list arrangement as the `Link` interface in Self-Referential Interfaces: each instance of a `Link` class contains a `next` member that points at the next link in the chain; the final link's `next` member contains a null value. So, what looks like a class including itself really expresses pointer semantics: the `next` data member contains a pointer to the next link in the chain.

You may be wondering at this point what the difference is then between the `Link` interface in Self-Referential Interfaces and the `Link` class shown above. The difference is that classes have *value* semantics, whereas proxies have *reference* semantics. To illustrate this, consider the `Link` *interface* from Self-Referential Interfaces once more:

Slice

```
interface Link {
    idempotent SomeType getValue();
    idempotent Link*    next();
};
```

Here, `getValue` and `next` are both operations and the return value of `next` is `Link*`, that is, next returns a *proxy*. A proxy has *reference* semantics, that is, it denotes an object somewhere. If you invoke the `getValue` operation on a `Link` proxy, a message is sent to the (possibly remote) servant for that proxy. In other words, for proxies, the object stays put in its server process and we access the state of the object via remote procedure calls. Compare this with the definition of our `Link` *class*:

Slice

```
class Link {
    SomeType value;
    Link next;
};
```

Here, `value` and `next` are data members and the type of next is `Link`, which has *value* semantics. In particular, while `next` looks and feels like a pointer, *it cannot denote an instance in a different address space*. This means that if we have a chain of `Link` instances, all of the instances are in our local address space and, when we read or write a value data member, we are performing local address space operations. This means that an operation that returns a `Link` instance, such as `getHead`, does not just return the head of the chain, *but the entire chain*, as shown:

Copyright 2016, ZeroC, Inc.

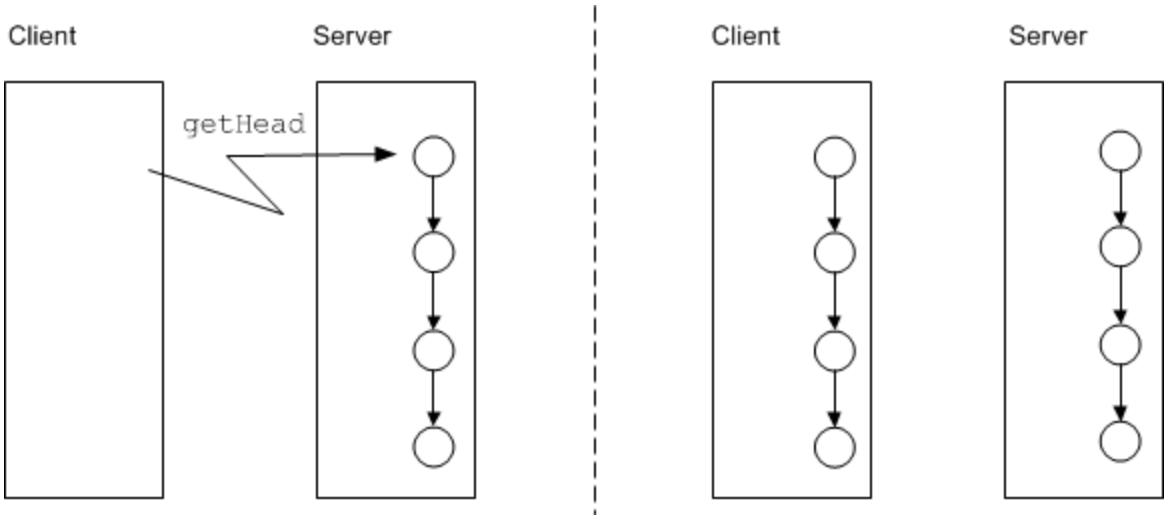

Class version of `Link` before and after calling `getHead`.

On the other hand, for the interface version of `Link`, we do not know where all the links are physically implemented. For example, a chain of four links could have each object instance in its own physical server process; those server processes could be each in a different continent. If you have a proxy to the head of this four-link chain and traverse the chain by invoking the `next` operation on each link, you will be sending four remote procedure calls, one to each object.

Self-referential classes are particularly useful to model graphs. For example, we can create a simple expression tree along the following lines:

Slice

```
enum UnaryOp { UnaryPlus, UnaryMinus, Not };
enum BinaryOp { Plus, Minus, Multiply, Divide, And, Or };

class Node {};

class UnaryOperator extends Node {
    UnaryOp operator;
    Node operand;
};

class BinaryOperator extends Node {
    BinaryOp op;
    Node operand1;
    Node operand2;
};

class Operand extends Node {
    long val;
};
```

The expression tree consists of leaf nodes of type `Operand`, and interior nodes of type `UnaryOperator` and `BinaryOperator`, with one or two descendants, respectively. All three of these classes are derived from a common base class `Node`. Note that `Node` is an empty class. This is one of the few cases where an empty base class is justified. (See the discussion on empty interfaces; once we add operations to this class hierarchy, the base class is no longer empty.)

If we write an operation that, for example, accepts a `Node` parameter, passing that parameter results in transmission of the entire tree to the server:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Evaluator {
    long eval(Node expression); // Send entire tree for evaluation
};
```

Self-referential classes are not limited to acyclic graphs; the Ice run time permits loops: it ensures that no resources are leaked and that infinite loops are avoided during marshaling.

See Also

- Classes with Operations
- Self-Referential Interfaces

Copyright 2016, ZeroC, Inc.

Classes Versus Structures

One obvious question to ask is: why does Ice provide structures as well as classes, when classes obviously can be used to model structures? The answer has to do with the cost of implementation: classes provide a number of features that are absent for structures:

- Classes support inheritance.
- Classes can be self-referential.
- Classes can have optional data members.
- Classes can have operations.
- Classes can implement interfaces.

Obviously, an implementation cost is associated with the additional features of classes, both in terms of the size of the generated code and the amount of memory and CPU cycles consumed at run time. On the other hand, structures are simple collections of values ("plain old structs") and are implemented using very efficient mechanisms. This means that, if you use structures, you can expect better performance and smaller memory footprint than if you would use classes (especially for languages with direct support for "plain old structures", such as C++ and C#). Use a class only if you need at least one of its more powerful features.

See Also

- Structures
- Classes with Operations
- Classes Implementing Interfaces

Copyright 2016, ZeroC, Inc.

Classes with Operations

Classes, in addition to data members, can have operations. The syntax for operation definitions in classes is identical to the syntax for operations in interfaces. For example, we can modify the expression tree from Self-Referential Classes as follows:

Slice

```
enum UnaryOp { UnaryPlus, UnaryMinus, Not };
enum BinaryOp { Plus, Minus, Multiply, Divide, And, Or };

class Node {
    idempotent long eval();
};

class UnaryOperator extends Node {
    UnaryOp operator;
    Node operand;
};

class BinaryOperator extends Node {
    BinaryOp op;
    Node operand1;
    Node operand2;
};

class Operand {
    long val;
};
```

The only change compared to the version in Self-Referential Classes is that the `Node` class now has an `eval` operation. The semantics of this are as for a virtual member function in C++: each derived class inherits the operation from its base class and can choose to override the operation's definition. For our expression tree, the `Operand` class provides an implementation that simply returns the value of its `val` member, and the `UnaryOperator` and `BinaryOperator` classes provide implementations that compute the value of their respective subtrees. If we call `eval` on the root node of an expression tree, it returns the value of that tree, regardless of whether we have a complex expression or a tree that consists of only a single `Operand` node.

Operations on classes are normally executed in the caller's address space, that is, operations on classes are *local* operations that do not result in a remote procedure call.

It is also possible to invoke an operation on a remote class instance.

Of course, this immediately raises an interesting question: what happens if a client receives a class instance with operations from a server, but client and server are implemented in different languages? Classes with operations require the receiver to supply a factory for instances of the class. The Ice run time only marshals the data members of the class. If a class has operations, the receiver of the class must provide a class factory that can instantiate the class in the receiver's address space, and the receiver is responsible for providing an implementation of the class's operations.

Therefore, if you use classes with operations, it is understood that client and server each have access to an implementation of the class's operations. No code is shipped over the wire (which, in an environment of heterogeneous nodes using different operating systems and languages is infeasible).

See Also

- Self-Referential Classes
- Pass-by-Value Versus Pass-by-Reference

Copyright 2016, ZeroC, Inc.

Architectural Implications of Classes

Classes have a number of architectural implications that are worth exploring in some detail.

On this page:

- Classes without Operations
- Classes with Operations
- Classes for Persistence

Classes without Operations

Classes that do not use inheritance and only have data members (whether self-referential or not) pose no architectural problems: they simply are values that are marshaled like any other value, such as a sequence, structure, or dictionary. Classes using derivation also pose no problems: if the receiver of a derived instance has knowledge of the derived type, it simply receives the derived type; otherwise, the instance is sliced to the most-derived type that is understood by the receiver. This makes class inheritance useful as a system is extended over time: you can create derived class without having to upgrade all parts of the system at once.

Classes with Operations

Classes with operations require additional thought. Here is an example: suppose that you are creating an Ice application. Also assume that the Slice definitions use quite a few classes with operations. You sell your clients and servers (both written in Java) and end up with thousands of deployed systems.

As time passes and requirements change, you notice a demand for clients written in C++.

For commercial reasons, you would like to leave the development of C++ clients to customers or a third party but, at this point, you discover a glitch: your application has lots of classes with operations along the following lines:

Slice

```
class ComplexThingForExpertsOnly {
    // Lots of arcane data members here...
    MysteriousThing mysteriousOperation(/* parameters */);
    ArcaneThing arcaneOperation(/* parameters */);
    ComplexThing complexOperation(/* parameters */);
    // etc...
};
```

It does not matter what exactly these operations do. (Presumably, you decided to off-load some of the processing for your application onto the client side for performance reasons.) Now that you would like other developers to write C++ clients, it turns out that your application will work only if these developers provide implementations of all the client-side operations and, moreover, if the semantics of these operations exactly match the semantics of your Java implementations. Depending on what these operations do, providing exact semantic equivalents in a different language may not be trivial, so you decide to supply the C++ implementations yourself.

But now, you discover another problem: the C++ clients need to be supported for a variety of operating systems that use a variety of different C++ compilers. Suddenly, your task has become quite daunting: you really need to supply implementations for all the combinations of operating systems and compiler versions that are used by clients. Given the different state of compliance with the ISO C++ standard of the various compilers, and the idiosyncrasies of different operating systems, you may find yourself facing a development task that is much larger than anticipated. And, of course, the same scenario will arise again should you need client implementations in yet another language.

The moral of this story is not that classes with operations should be avoided; they can provide significant performance gains and are not necessarily bad. But, keep in mind that, once you use classes with operations, you are, in effect, using client-side native code and, therefore, you can no longer enjoy the implementation transparencies that are provided by interfaces. This means that classes with operations should be used only if you can tightly control the deployment environment of clients. If not, you are better off using interfaces and classes without operations. That way, all the processing stays on the server and the contract between client and server is provided solely by the Slice definitions, not by the semantics of the additional client-side code that is required for classes with operations.

Classes for Persistence

Ice also provides a built-in persistence mechanism that allows you to store the state of a class in a database with very little implementation effort. To get access to these persistence features, you must define a Slice class whose members store the state of the class.

Copyright 2016, ZeroC, Inc.

See Also

- Freeze

Copyright 2016, ZeroC, Inc.

Classes Implementing Interfaces

A Slice class can also be used as a servant in a server, that is, an instance of a class can be used to provide the behavior for an interface, for example:

Slice

```
interface Time {
    idempotent TimeOfDay getTime();
    idempotent void setTime(TimeOfDay time);
};

class Clock implements Time {
    TimeOfDay time;
};
```

The `implements` keyword indicates that the class `Clock` provides an *implementation* of the `Time` interface. The class can provide data members and operations of its own; in the preceding example, the `Clock` class stores the current time that is accessed via the `Time` interface. A class can implement several interfaces, for example:

Slice

```
interface Time {
    idempotent TimeOfDay getTime();
    idempotent void setTime(TimeOfDay time);
};

interface Radio {
    idempotent void setFrequency(long hertz);
    idempotent void setVolume(long dB);
};

class RadioClock implements Time, Radio {
    TimeOfDay time;
    long hertz;
};
```

The class `RadioClock` implements both `Time` and `Radio` interfaces.

A class, in addition to implementing an interface, can also extend another class:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Time {
    idempotent TimeOfDay getTime();
    idempotent void setTime(TimeOfDay time);
};

class Clock implements Time {
    TimeOfDay time;
};

interface AlarmClock extends Time {
    idempotent TimeOfDay getAlarmTime();
    idempotent void setAlarmTime(TimeOfDay alarmTime);
};

interface Radio {
    idempotent void setFrequency(long hertz);
    idempotent void setVolume(long dB);
};

class RadioAlarmClock extends Clock
                      implements AlarmClock, Radio {
    TimeOfDay alarmTime;
    long hertz;
};
```

These definitions result in the following inheritance graph:

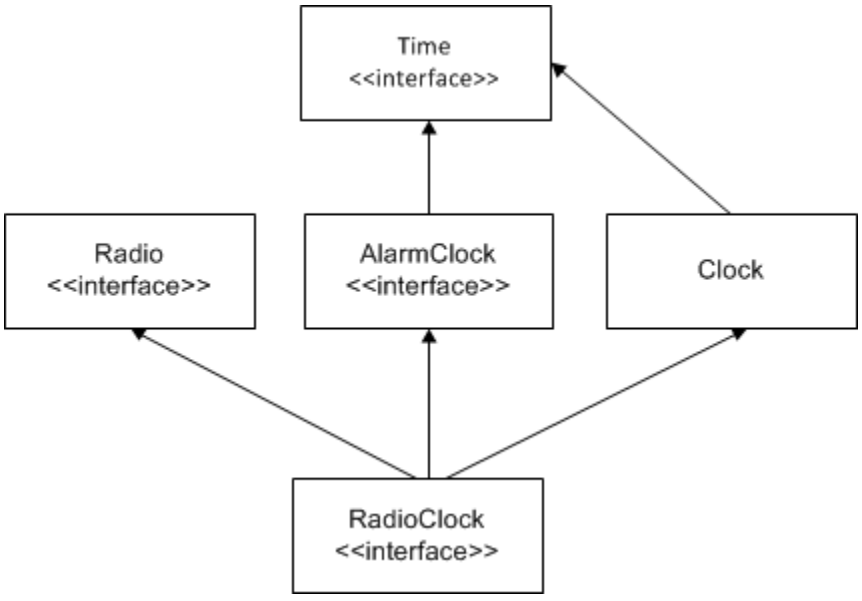

A Class using implementation and interface inheritance.

For this definition, `Radio` and `AlarmClock` are abstract interfaces, and `Clock` and `RadioAlarmClock` are concrete classes. As for Java, a class can implement multiple interfaces, but can extend at most one class.

Copyright 2016, ZeroC, Inc.

See Also

- Architectural Implications of Classes
- Class Inheritance Limitations

Copyright 2016, ZeroC, Inc.

Class Inheritance Limitations

As for interface inheritance, a class cannot redefine an operation or data member that it inherits from a base interface or class. For example:

Slice

```
interface BaseInterface {
    void op();
};

class BaseClass {
    int member;
};

class DerivedClass extends BaseClass implements BaseInterface {
    void someOperation();       // OK
    int op();                   // Error!
    int  someMember;            // OK
    long member;                // Error!
};
```

See Also

- Interface Inheritance

Copyright 2016, ZeroC, Inc.

Pass-by-Value Versus Pass-by-Reference

As we saw in Self-Referential Classes, classes naturally support pass-by-value semantics: passing a class transmits the data members of the class to the receiver. Any changes made to these data members by the receiver affect only the receiver's copy of the class; the data members of the sender's class are not affected by the changes made by the receiver.

In addition to passing a class by value, you can pass a class by reference. For example:

Slice

```
class TimeOfDay {
    short hour;
    short minute;
    short second;
    string format();
};

interface Example {
     TimeOfDay* get();  // Note: returns a proxy!
};
```

Note that the `get` operation returns a *proxy* to a `TimeOfDay` class and not a `TimeOfDay` instance itself. The semantics of this are as follows:

- When the client receives a `TimeOfDay` proxy from the `get` call, it holds a proxy that differs in no way from an ordinary proxy for an interface.
- The client can invoke operations via the proxy, but *cannot* access the data members. This is because proxies do not have the concept of data members, but represent interfaces: even though the `TimeOfDay` class has data members, only its *operations* can be accessed via a the proxy.

The net effect is that, in the preceding example, the server holds an instance of the `TimeOfDay` class. A proxy for that instance was passed to the client. The only thing the client can do with this proxy is to invoke the `format` operation. The implementation of that operation is provided by the server and, when the client invokes `format`, it sends an RPC message to the server just as it does when it invokes an operation on an interface. The implementation of the `format` operation is entirely up to the server. (Presumably, the server will use the data members of the `TimeOfDay` instance it holds to return a string containing the time to the client.)

The preceding example looks somewhat contrived for classes only. However, it makes perfect sense if classes implement interfaces: parts of your application can exchange class instances (and, therefore, state) by value, whereas other parts of the system can treat these instances as remote interfaces.

For example:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Time {
    string format();
    // ...
};

class TimeOfDay implements Time {
    short hour;
    short minute;
    short second;
};

interface I1 {
     TimeOfDay get();           // Pass by value
     void put(TimeOfDay time);  // Pass by value
};

interface I2 {
    Time* get();                // Pass by reference
};
```

In this example, clients dealing with interface `I1` are aware of the `TimeOfDay` class and pass it by value whereas clients dealing with interface `I2` deal only with the `Time` interface. However, the actual implementation of the `Time` interface in the server uses `TimeOfDay` inst ances.

Be careful when designing systems that use such mixed pass-by-value and pass-by-reference semantics. Unless you are clear about what parts of the system deal with the interface (pass by reference) aspects and the class (pass by value) aspects, you can end up with something that is more confusing than helpful.

A good example of putting this feature to use can be found in Freeze, which allows you to add classes to an existing interface to implement persistence.

See Also

- Self-Referential Classes
- Freeze

Copyright 2016, ZeroC, Inc.

Passing Interfaces by Value

Consider the following definitions:

Slice

```
interface Time {
    idempotent TimeOfDay getTime();
    // ...
};

interface Record {
    void addTimeStamp(Time t); // Note: Time t, not Time* t
    // ...
};
```

Note that `addTimeStamp` accepts a parameter of type `Time`, not of type `Time*`. The question is, what does it mean to pass an interface *by value*? Obviously, at run time, we cannot pass an an actual interface to this operation because interfaces are abstract and cannot be instantiated. Neither can we pass a proxy to a `Time` object to `addTimeStamp` because a proxy cannot be passed where an interface is expected.

However, what we *can* pass to `addTimeStamp` is something that is not abstract and derives from the `Time` interface. For example, at run time, we could pass an instance of the `TimeOfDay` class we saw earlier. Because the `TimeOfDay` class derives from the `Time` interface, the class type is compatible with the formal parameter type `Time` and, at run time, what is sent over the wire to the server is the `TimeOfDay` class instance.

See Also

- Pass-by-Value Versus Pass-by-Reference

Copyright 2016, ZeroC, Inc.

Classes with Compact Type IDs

As of Ice 3.5, you can optionally associate a numeric identifier with a class. The Ice run time substitutes this value, known as a *compact type ID*, in place of its equivalent string type ID during marshaling to conserve space. The compact type ID follows immediately after the class name, enclosed in parentheses:

Slice

```
module MyModule {
class CompactExample(4) {
    // ...
};
};
```

In this example, the Ice run time marshals the value 4 instead of its string equivalent `"::MyModule::CompactExample"`. The specified value must be a non-negative integer that is unique within the translation unit.

Using values less than 255 produces the most efficient encoding.

See Also

- Classes
- Type IDs
- Data Encoding for Class Type IDs

Copyright 2016, ZeroC, Inc.

Forward Declarations

Both interfaces and classes can be forward declared. Forward declarations permit the creation of mutually dependent objects, for example:

Slice

```
module Family {
    interface Child;            // Forward declaration

    sequence<Child*> Children;  // OK

    interface Parent {
        Children getChildren(); // OK
    };

    interface Child {           // Definition
        Parent* getMother();
        Parent* getFather();
    };
};
```

Without the forward declaration of `Child`, the definition obviously could not compile because `Child` and `Parent` are mutually dependent interfaces. You can use forward-declared interfaces and classes to define types (such as the `Children` sequence in the previous example). Forward-declared interfaces and classes are also legal as the type of a structure, exception, or class member, as the value type of a dictionary, and as the parameter and return type of an operation. However, you cannot inherit from a forward-declared interface or class until after its definition has been seen by the compiler:

Slice

```
interface Base;                       // Forward declaration

interface Derived1 extends Base {};   // Error!

interface Base {};                    // Definition

interface Derived2 extends Base {};   // OK, definition was seen
```

Not inheriting from a forward-declared base interface or class until its definition is seen is necessary because, otherwise, the compiler could not enforce that derived interfaces must not redefine operations that appear in base interfaces.

A multi-pass compiler could be used, but the added complexity is not worth it.

See Also

- Interfaces, Operations, and Exceptions
- Classes

Copyright 2016, ZeroC, Inc.

Optional Data Members

On this page:

- Overview of Optional Data Members
- Declaring Optional Data Members
- Optional Data Members with Default Values

Overview of Optional Data Members

As of Ice 3.5, a data member of a Slice class or exception may be declared as optional to indicate that a program can leave its value unset. Data members not declared as optional are known as *required* members; a program must supply legal values for all required members.

Declaring Optional Data Members

Each optional data member in a type must be assigned a unique, non-negative integer *tag*:

Slice

```
class C
{
    string name;
    bool active;
    optional(2) string alternateName;
    optional(5) int overrideCode;
};
```

It is legal for a base type's tag to be reused by a derived type:

Slice

```
exception Base
{
    optional(1) int systemCode;
};

exception Derived extends Base
{
    optional(1) string diagnostic; // OK
};
```

The scope of a tag is limited to its enclosing type and has no effect on base or derived types.

Language mappings specify an API for setting an optional member and testing whether a member is set. Here is an example in C++:

Copyright 2016, ZeroC, Inc.

C++

```
CPtr c = new C;
c->name = "xyz";           // required
c->active = true;          // required
c->alternateName = "abc";  // optional
c->overrideCode = 42;      // optional

if(c->alternateName)
    cout << "alt name = " << c->alternateName << endl;
```

As you can see, the C++ language mapping makes setting an optional member as simple as assigning it a value. Refer to the language mapping sections for more details on the optional data member API.

A well-behaved program must test for the presence of an optional member and not assume that it is always set. Dereferencing an unset optional member causes a run-time error.

In all supported language mappings, an optional data member's initial condition is unset if not otherwise assigned during construction. Again using C++ as an example:

C++

```
CPtr c = new C;            // default constructor
assert(!c->alternateName); // not set

c = new C("xyz", true, "abc", 42); // one-shot constructor
assert(c->alternateName);          // set by constructor
```

Optional Data Members with Default Values

You can declare a default value for optional members just as you can for required members:

Slice

```
class C
{
    string name;
    bool active = true;
    optional(2) string alternateName;
    optional(5) int overrideCode = -1;
};
```

An optional data member with a default value is considered to be set by default:

Copyright 2016, ZeroC, Inc.

C++

```
CPtr c = new C;            // default constructor
assert(!c->alternateName); // not set
assert(c->overrideCode);   // set to default value
```

Each language mapping provides an API for resetting an optional data member to its unset condition.

See Also

- Classes
- User Exceptions
- Optional Values

Copyright 2016, ZeroC, Inc.

Type IDs

Interface, class and user exception Slice types have an internal type identifier, known as the *type ID*. The type ID is simply the fully-qualified name of each type. For example, the type ID of the `Child` interface in the preceding example is `::Family::Children::Child`. A type ID starts with a leading `::` and is formed by starting with the global scope (`::`) and forming the fully-qualified name of a type by appending each module name in which the type is nested, and ending with the name of the type itself; the components of the type ID are separated by `::`.

The type ID of the Slice `Object` type is `::Ice::Object`.

Type IDs are used internally by the Ice run time as a unique identifier for each type. For example, when an exception is raised, the marshaled form of the exception that is returned to the client is preceded by its type ID on the wire. The client-side run time first reads the type ID and, based on that, unmarshals the remainder of the data as appropriate for the type of the exception.

Type IDs are also used by the `ice_isA` operation.

See Also

- ice_isA

Copyright 2016, ZeroC, Inc.

Operations on Object

The `Object` interface has a number of operations. We cannot define type `Object` in Slice because `Object` is a keyword; regardless, here is what (part of) the definition of `Object` would look like if it were legal:

Slice

```
sequence<string> StrSeq;

interface Object {                      // "Pseudo" Slice!
    idempotent void   ice_ping();
    idempotent bool   ice_isA(string typeID);
    idempotent string ice_id();
    idempotent StrSeq ice_ids();
    // ...
};
```

Note that, apart from the illegal use of the keyword `Object` as the interface name, the operation names all contain the `ice_` prefix. This prefix is reserved for use by Ice and cannot clash with a user-defined operation. This means that all Slice interfaces can inherit from `Object` without name clashes. We discuss these built-in operations below.

On this page:

- ice_ping
- ice_isA
- ice_id
- ice_ids

ice_ping

All interfaces support the `ice_ping` operation. That operation is useful for debugging because it provides a basic reachability test for an object: if the object exists and a message can successfully be dispatched to the object, `ice_ping` simply returns without error. If the object cannot be reached or does not exist, `ice_ping` throws a run-time exception that provides the reason for the failure.

ice_isA

The `ice_isA` operation accepts a type identifier (such as the identifier returned by `ice_id`) and tests whether the target object supports the specified type, returning `true` if it does. You can use this operation to check whether a target object supports a particular type. For example, referring to the diagram Implicit Inheritance from Object once more, assume that you are holding a proxy to a target object of type `AlarmClock`. The table below illustrates the result of calling `ice_isA` on that proxy with various arguments. (We assume that all types in the Implicit inheritance from Object diagram are defined in a module `Times`):

Argument	Result
`::Ice::Object`	`true`
`::Times::Clock`	`true`
`::Times::AlarmClock`	`true`
`::Times::Radio`	`false`
`::Times::RadioClock`	`false`

Calling `ice_isA` on a proxy denoting an object of type AlarmClock.

As expected, `ice_isA` returns true for `::Times::Clock` and `::Times::AlarmClock` and also returns true for `::Ice::Object` (because all interfaces support that type). Obviously, an `AlarmClock` supports neither the `Radio` nor the `RadioClock` interfaces, so `ice_isA` returns false for these types.

Copyright 2016, ZeroC, Inc.

ice_id

The `ice_id` operation returns the type ID of the most-derived type of an interface.

ice_ids

The `ice_ids` operation returns a sequence of type IDs that contains all of the type IDs supported by an interface. For example, for the RadioClock interface in Implicit inheritance from Object, `ice_ids` returns a sequence containing the type IDs `::Ice::Object`, `::Times::Clock`, `::Times::AlarmClock`, `::Times::Radio`, and `::Times::RadioClock`.

See Also

- Type IDs
- Interface Inheritance
- Implicit inheritance from Object

Copyright 2016, ZeroC, Inc.

Local Types

In order to access certain features of the Ice run time, you must use APIs that are provided by libraries. However, instead of defining an API that is specific to each implementation language, Ice defines its APIs in Slice using the `local` keyword. The advantage of defining APIs in Slice is that a single definition suffices to define the API for all possible implementation languages. The actual language-specific API is then generated by the Slice compiler for each implementation language. Types that are provided by Ice libraries are defined using the Slice `local` keyword.

For example:

Slice

```
module Ice {
    local interface ObjectAdapter {
        // ...
    };
};
```

Any Slice definition (not just interfaces) can have a `local` modifier. If the `local` modifier is present, the Slice compiler does not generate marshaling code for the corresponding type. This means that a local type can *never* be accessed remotely because it cannot be transmitted between client and server. (The Slice compiler prevents use of `local` types in non-`local` contexts.)

In addition, local interfaces and local classes do *not* inherit from `Ice::Object`. Instead, local interfaces and classes have their own, completely separate inheritance hierarchy. At the root of this hierarchy is the type `Ice::LocalObject`, as shown:

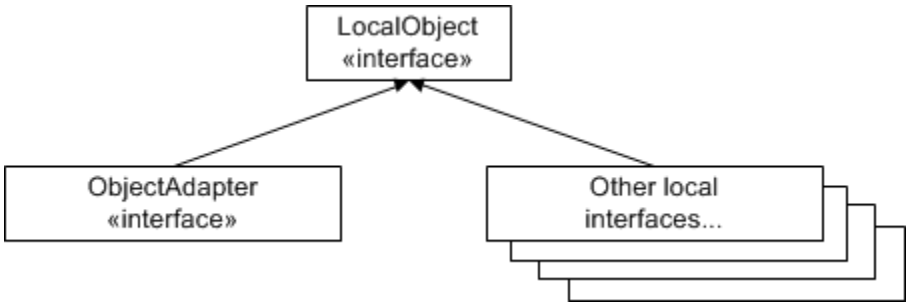

Inheritance from `LocalObject`.

Because local interfaces form a completely separate inheritance hierarchy, you cannot pass a local interface where a non-local interface is expected, and vice-versa.

You rarely need to define local types for your own applications — the `local` keyword exists mainly to allow definition of APIs for the Ice run time. (Because local objects cannot be invoked remotely, there is little point for an application to define local objects; it might as well define ordinary programming-language objects instead.) However, there is one exception to this rule: servant locators must be implemented as local objects.

See Also

- Servant Locators

Copyright 2016, ZeroC, Inc.

Names and Scoping

Slice has a number of rules regarding identifiers. You will typically not have to concern yourself with these. However, occasionally, it is good to know how Slice uses naming scopes and resolves identifiers.

On this page:

- Naming Scope
- Case Sensitivity
- Qualified Names
- Names in Nested Scopes
- Introduced Identifiers
- Name Lookup Rules
- Scoping Rules for Parameters and Data Members
- Scoping Rules in Prior Ice Releases

Naming Scope

The following Slice constructs establish a naming scope:

- the global (file) scope
- modules
- interfaces
- classes
- structures
- exceptions
- parameter lists

Within a naming scope, identifiers must be unique, that is, you cannot use the same identifier for different purposes. For example:

Slice

```
interface Bad {
    void op(int p, string p);   // Error!
};
```

Because a parameter list forms a naming scope, it is illegal to use the same identifier p for different parameters. Similarly, data members, operation names, interface and class names, etc. must be unique within their enclosing scope.

Case Sensitivity

Identifiers that differ only in case are considered identical, so you must use identifiers that differ not only in capitalization within a naming scope. For example:

Slice

```
struct Bad {
    int    m;
    string M;   // Error!
};
```

The Slice compiler also enforces consistent capitalization for identifiers. Once you have defined an identifier, you must use the same capitalization for that identifier thereafter. For example, the following is in error:

Copyright 2016, ZeroC, Inc.

Slice

```
sequence<string> StringSeq;

interface Bad {
    stringSeq op();     // Error!
};
```

Note that identifiers must not differ from a Slice keyword in case only. For example, the following is in error:

Slice

```
interface Module {      // Error, "module" is a keyword
    // ...
};
```

Qualified Names

The scope-qualification operator `::` allows you to refer to a type in a non-local scope. For example:

Slice

```
module Types {
    sequence<long> LongSeq;
};

module MyApp {
    sequence<Types::LongSeq> NumberTree;
};
```

Here, the qualified name `Types::LongSeq` refers to `LongSeq` defined in module `Types`. The global scope is denoted by a leading `::`, so we could also refer to `LongSeq` as `::Types::LongSeq`.

The scope-qualification operator also allows you to create mutually dependent interfaces that are defined in different modules. The obvious attempt to do this fails:

Copyright 2016, ZeroC, Inc.

Slice

```
module Parents {
    interface Children::Child;  // Syntax error!
    interface Mother {
        Children::Child* getChild();
    };
    interface Father {
        Children::Child* getChild();
    };
};

module Children {
    interface Child {
        Parents::Mother* getMother();
        Parents::Father* getFather();
    };
};
```

This fails because it is syntactically illegal to forward-declare an interface in a different module. To make it work, we must use a reopened module:

Slice

```
module Children {
    interface Child;                    // Forward declaration
};

module Parents {
    interface Mother {
        Children::Child* getChild();    // OK
    };
    interface Father {
        Children::Child* getChild();    // OK
    };
};

module Children {                       // Reopen module
    interface Child {                   // Define Child
        Parents::Mother* getMother();
        Parents::Father* getFather();
    };
};
```

While this technique works, it is probably of dubious value: mutually dependent interfaces are, by definition, tightly coupled. On the other hand, modules are meant to be used to place related definitions into the same module, and unrelated definitions into different modules. Of course, this begs the question: if the interfaces are so closely related that they depend on each other, why are they defined in different modules? In the interest of clarity, you probably should avoid this construct, even though it is legal.

Copyright 2016, ZeroC, Inc.

Names in Nested Scopes

Names defined in an enclosing scope can be redefined in an inner scope. For example, the following is legal:

Slice

```
module Outer {
    sequence<string> Seq;

    module Inner {
        sequence<short> Seq;
    };
};
```

Within module `Inner`, the name `Seq` refers to a sequence of `short` values and hides the definition of `Outer::Seq`. You can still refer to the other definition by using explicit scope qualification, for example:

Slice

```
module Outer {
    sequence<string> Seq;

    module Inner {
        sequence<short> Seq;

        struct Confusing {
            Seq          a;     // Sequence of short
            ::Outer::Seq b;     // Sequence of string
        };
    };
};
```

Needless to say, you should try to avoid such redefinitions — they make it harder for the reader to follow the meaning of a specification.

Same-named constructs cannot be nested inside each other in certain situations. For example, a module named `M` cannot (recursively) contain any construct also named `M`. The same is true for interfaces and classes, which cannot define an operation with the same name as the enclosing interface or class. For example, the following examples are all in error:

Copyright 2016, ZeroC, Inc.

Slice

```
module M {
    interface M { /* ... */ };  // Error!

    interface I {
        void I();               // Error!
    };
};

module Outer {
    module Inner {
        interface outer {       // Error, even if case differs!
            // ...
        };
    };
};
```

The reason for this restriction is that nested types that have the same name are difficult to map into some languages. For example, C++ and Java reserve the name of a class as the name of the constructor, so an interface `I` could not contain an operation named `I` without artificial rules to avoid the name clash.

Similarly, some languages (such as C# prior to version 2.0) do not permit a qualified name to be anchored at the global scope. If a nested module or type is permitted to have the same name as the name of an enclosing module, it can become impossible to generate legal code in some cases.

In the interest of simplicity, Slice prohibits the name of a nested module or type from being the same as the name of one of its enclosing modules.

Introduced Identifiers

Within a naming scope, an identifier is introduced at the point of first use; thereafter, within that naming scope, the identifier cannot change meaning.

For example:

Slice

```
module M {
    sequence<string> Seq;

    interface Bad {
        Seq op1();      // Seq and op1 introduced here
        int Seq();      // Error, Seq has changed meaning
    };
};
```

The declaration of `op1` uses `Seq` as its return type, thereby introducing `Seq` into the scope of interface `Bad`. Thereafter, `Seq` can only be used as a type name that denotes a sequence of strings, so the compiler flags the declaration of the second operation as an error.

Note that fully-qualified identifiers are not introduced into the current scope:

Copyright 2016, ZeroC, Inc.

Slice

```
module M {
    sequence<string> Seq;

    interface Bad {
        ::M::Seq op1(); // Only op1 introduced here
        int Seq();      // OK
    };
};
```

In general, a fully-qualified name (one that is anchored at the global scope and, therefore, begins with a `::` scope resolution operator) does not introduce any name into the current scope. On the other hand, a qualified name that is not anchored at the global scope introduces only the first component of the name:

Slice

```
module M {
    sequence<string> Seq;

    interface Bad {
        M::Seq op1();   // M and op1 introduced here, but not Seq
        int Seq();      // OK
    };
};
```

Name Lookup Rules

When searching for the definition of a name that is not anchored at the global scope, the compiler first searches backward in the current scope of a definition of the name. If it can find the name in the current scope, it uses that definition. Otherwise, the compiler successively searches enclosing scopes for the name until it reaches the global scope. Here is an example to illustrate this:

Copyright 2016, ZeroC, Inc.

Slice

```
module M1 {
    sequence<double> Seq;

    module M2 {
        sequence<string> Seq;   // OK, hides ::M1::Seq

        interface Base {
            Seq op1();          // Returns sequence of string
        };
    };

    module M3 {
        interface Derived extends M2::Base {
            Seq op2();          // Returns sequence of double
        };

        sequence<bool> Seq;     // OK, hides ::M1::Seq

        interface I {
            Seq op();           // Returns sequence of bool
        };
    };

    interface I {
        Seq op();               // Returns sequence of double
    };
};
```

Note that `M2::Derived::op2` returns a sequence of `double`, even though `M1::Base::op1` returns a sequence of `string`. That is, the meaning of a type in a base interface is irrelevant to determining its meaning in a derived interface — the compiler always searches for a definition only in the current scope and enclosing scopes, and never takes the meaning of a name from a base interface or class.

Scoping Rules for Parameters and Data Members

A Slice operation creates a new naming scope in which all parameter names must be unique:

Slice

```
interface I {
    void op1(string p, int P);    // Error, differs only in case
    void op2(int n, out int n);   // Error, duplicate
    void op3(string s, int i);    // OK
};
```

It's legal for parameters to reuse the names of symbols in enclosing scopes, including the name of the operation, class, interface or module:

Copyright 2016, ZeroC, Inc.

Slice

```
module M {
    sequence<string> Seq;

    interface I {
        string query(string query);   // OK to reuse operation name
        void op1(int I);              // OK to reuse name of enclosing type
        void op2(Seq Seq);            // OK to reuse type name
        void op3(int M);              // OK to reuse module name
    };
};
```

The rules for data members are similar to those of parameters:

- Structures
  Member names must be unique within the structure.
- Exceptions
  Member names must be unique within the exception, including any members inherited from base exceptions.
- Classes
  Member names must be unique within the class, including any members inherited from base classes. Members must not duplicate the names of operations defined by the class or inherited by any base classes or interfaces.

As for parameters, data members can reuse the names of symbols in enclosing scopes. The examples below illustrate these rules:

Copyright 2016, ZeroC, Inc.

Slice

```
module M {
    struct S {
        int i;
        string s;    // OK to reuse name of enclosing type
        long I;      // Error, differs only in case
        bool M;      // OK to reuse module name
    };

    interface I {
        void op();
    };

    class Base {
        string name;
    };

    class C extends Base implements I {
        S S;         // OK to reuse type name
        byte c;      // OK to reuse name of enclosing type
        string op;   // Error, duplicates inherited I::op
        string Name; // Error, differs only in case from Base::name
    };

    exception ErrorBase {
        string reason;
    };

    exception Error extends ErrorBase {
        long error;  // OK to reuse name of enclosing type
        int reason;  // Error, duplicates inherited ErrorBase::reason
    };
};
```

Scoping Rules in Prior Ice Releases

The scoping rules for parameters and data members were more restrictive in Ice 3.5 and earlier releases:

- A data member cannot have the same name as its enclosing type

  Slice

  ```
  class C {
      int c; // Error
  };
  ```

- A data member cannot have the same name as its type

Copyright 2016, ZeroC, Inc.

Slice

```
module M {
    sequence<string> Seq;
    struct S {
        Seq Seq; // Error, use ::M::Seq as the type instead
    };
};
```

You can work around this limitation by using the fully-qualified type name.

- A parameter cannot have the same name as its operation

Slice

```
void op(int op); // Error
```

- A parameter cannot have the same name as its type

Slice

```
module M {
    sequence<string> Seq;
    interface I {
        void op(Seq Seq); // Error, use ::M::Seq as the type instead
    };
};
```

You can work around this limitation by using the fully-qualified type name.

⚠ You must adhere to these more-restrictive rules if your Slice definitions need to maintain backward compatibility with Ice 3.5 or earlier.

See Also

- Lexical Rules

Copyright 2016, ZeroC, Inc.

Metadata

Slice has the concept of a *metadata* directive. For example:

Slice

```
["java:type:java.util.LinkedList<Integer>"] sequence<int> IntSeq;
```

A metadata directive can appear as a prefix to any Slice definition. Metadata directives appear in a pair of square brackets and contain one or more string literals separated by commas. For example, the following is a syntactically valid metadata directive containing two strings:

Slice

```
["a", "b"] interface Example {};
```

Metadata directives are not part of the Slice language per se: the presence of a metadata directive has no effect on the client-server contract, that is, metadata directives do not change the Slice type system in any way. Instead, metadata directives are targeted at specific back-ends, such as the code generator for a particular language mapping. In the preceding example, the `java:` prefix indicates that the directive is targeted at the Java code generator.

Metadata directives permit you to provide supplementary information that does not change the Slice types being defined, but somehow influences how the compiler will generate code for these definitions. For example, a metadata directive `java:type:java.util.LinkedList<T>` instructs the Java code generator to map a sequence to a linked list instead of an array (which is the default).

Metadata directives are also used to create skeletons that support *Asynchronous Method Dispatch (AMD)*.

Apart from metadata directives that are attached to a specific definition, there are also global metadata directives. For example:

Slice

```
[["java:package:com.acme"]]
```

Note that a global metadata directive is enclosed by double square brackets, whereas a local metadata directive (one that is attached to a specific definition) is enclosed by single square brackets. Global metadata directives are used to pass instructions that affect the entire compilation unit. For example, the preceding metadata directive instructs the Java code generator to generate the contents of the source file into the Java package `com.acme`. Global metadata directives must precede any definitions in a file (but can appear following any `#include` directives).

We discuss specific metadata directives in the relevant chapters to which they apply.

You can find a summary of all metadata directives in Slice Metadata Directives.

See Also

- Slice Metadata Directives

Copyright 2016, ZeroC, Inc.

Serializable Objects

Ice for Java and Ice for .NET allow you to send native Java and CLR objects as operation parameters. The Ice run time automatically serializes and deserializes the objects as part of an invocation. This mechanism allows you to transmit Java and CLR objects that do not have a corresponding Slice definition.

On this page:

- The serializable Metadata Directive
- Architectural Implications

The `serializable` Metadata Directive

To enable serialization, the parameter type must be a byte sequence with appropriate metadata. For example:

Slice

```
["java:serializable:SomePackage.JavaClass"]
sequence<byte> JavaObj;

interface JavaExample {
    void sendJavaObj(JavaObj o);
};

["clr:serializable:SomeNamespace.CLRClass"]
sequence<byte> CLRObj;

interface CLRExample {
    void sendCLRObj(CLRObj o);
};
```

The `java:serializable` metadata indicates that the corresponding byte sequence holds a Java serializable type named `SomePackage.JavaClass`. Your program must provide an implementation of this class; the class must implement `java.io.Serializable`.

Similarly, the `clr:serializable` metadata indicates that the corresponding byte sequences holds a CLR serializable type named `SomeNamespace.CLRClass`. Your program must provide an implementation of this class; the class must be marked with the `Serializable` attribute.

Architectural Implications

The `serializable` metadata directive permits you to transmit arbitrary Java and CLR objects across the network without the need to define corresponding Slice classes or structures. This is mainly a convenience feature: you could achieve the same thing by using ordinary Slice byte sequences and explicitly serializing your Java or CLR objects into byte sequences at the sending end, and deserializing them at the receiving end. The `serializable` metadata conveniently takes care of these chores for you and so is simpler to use.

Despite its convenience, you should use this feature with caution because it destroys language transparency. For example, a serialized Java object is useless to a C++ server. All the C++ server can do with such an object is to pass it on to some other process as a byte sequence. (Of course, if that receiving process is a Java process, it can deserialize the byte sequence.)

Further, similar to Slice classes with methods, a serialized object can be deserialized only if client and server agree on the definition of the serialized class. In Java, this is enforced by the `serialVersionUID` field of each instance; in the CLR, client and server must reference identical assembly versions. This creates much tighter coupling of client and server than exchanging Slice-defined types.

And, of course, if you build a system that relies on, for example, the exchange of serialized Java objects and you later find that you need to add C++ or C# components to the system, these components cannot do anything with the serialized Java objects other than pass them around as a blob of bytes.

So, if you do use these features, be clear that this implies tighter coupling between client and server, and that it creates additional library versioning and distribution issues because all parts of the system must agree on the implementation of the serialized objects.

Copyright 2016, ZeroC, Inc.

See Also

- Serializable Objects in Java
- Serializable Objects in C#
- Architectural Implications of Classes

Copyright 2016, ZeroC, Inc.

Deprecating Slice Definitions

All Slice compilers support a metadata directive that allows you to deprecate a Slice definition. For example:

Slice

```
interface Example {
    ["deprecate:someOperation() has been deprecated, use alternativeOperation() ins
tead."]
    void someOperation();

    void alternativeOperation();
};
```

The `["deprecate"]` metadata directive causes the compiler to emit code that generates a warning if you compile application code that uses a deprecated feature. This is useful if you want to remove a feature from a Slice definition but do not want to cause a hard error.

The message that follows the colon is optional; if you omit the message and use `["deprecate"]`, the Slice compilers insert a default message into the generated code.

You can apply the `["deprecate"]` metadata directive to Slice constructs other than operations (for example, a structure or sequence definition).

See Also

- Generating Slice Documentation

Copyright 2016, ZeroC, Inc.

Using the Slice Compilers

Ice provides a separate Slice compiler for each language mapping, as shown below:

Language	Compiler	Freeze Compiler
C++	`slice2cpp`	`slice2freeze`
Java	`slice2java`	`slice2freezej`
JavaScript	`slice2js`	
C#	`slice2cs`	
Objective-C	`slice2objc`	
Python	`slice2py`	
Ruby	`slice2rb`	
PHP	`slice2php`	

The Slice compilers.

The compilers share a similar command-line syntax:

```
<compiler-name> [options] file...
```

Regardless of which compiler you use, a number of command-line options are common to the compilers for any language mapping. (See the appropriate language mapping chapter for options that are specific to a particular language mapping.) The common command-line options are:

- `-h, --help`
  Displays a help message.
- `-v, --version`
  Displays the compiler version.
- `-D`*`NAME`*
  Defines the preprocessor symbol *`NAME`*.
- `-D`*`NAME=DEF`*
  Defines the preprocessor symbol *`NAME`* with the value *`DEF`*.
- `-U`*`NAME`*
  Undefines the preprocessor symbol {*`NAME`*.
- `-I`*`DIR`*
  Add the directory *`DIR`* to the search path for `#include` directives.
- `-E`
  Print the preprocessor output on `stdout`.
- `--output-dir` *`DIR`*
  Place the generated files into directory *`DIR`*.
- `-d, --debug`
  Print debug information showing the operation of the Slice parser.
- `--ice`
  Permit use of the reserved `Ice` prefix in Slice identifiers. This is useful when building Ice itself, or when compiling a Slice file that includes another Slice file with Ice-prefixed identifiers, such as `Ice/BuiltinSequences.ice`.
- `--underscore`
  Permit use of underscores in Slice identifiers.

Copyright 2016, ZeroC, Inc.

- `--depend`
  Print dependency information to standard output by default, or to the file specified by the `--depend-file` option. No code is generated when this option is specified. The output includes the complete list of Slice files that the input Slice files depend on through direct or indirect inclusion; this output may include other files depending on the target programming language. The Ice for C++ build system uses the script `config/makedepend.py` to process and include this output in `Makefiles`.

- `--depend-xml`
  Print dependency information in XML format to standard output by default, or to the file specified by the `--depend-file` option. No code is generated when this option is specified. The output consists of the complete list of Slice files that the input Slice files depend on through direct or indirect inclusion, and is identical will all Slice compilers.

- `--depend-file`
  Directs dependency information to the specified file. The output format depends on whether `--depend` or `--depend-xml` is also specified.

The Slice compilers permit you to compile more than a single source file, so you can compile several Slice definitions at once, for example:

```
slice2cpp -I. file1.ice file2.ice file3.ice
```

See Also

- Slice Compilation

Copyright 2016, ZeroC, Inc.

Slice Checksums

As distributed applications evolve, developers and system administrators must be careful to ensure that deployed components are using the same client-server contract. Unfortunately, mistakes do happen, and it is not always readily apparent when they do.

To minimize the chances of this situation, the Slice compilers support an option that generates checksums for Slice definitions, thereby enabling two peers to verify that they share an identical client-server contract. The checksum for a Slice definition includes details such as parameter and member names and the order in which operations are defined, but ignores information that is not relevant to the client-server contract, such as metadata, comments, and formatting.

This option causes the Slice compiler to construct a dictionary that maps Slice type identifiers to checksums. A server typically supplies an operation that returns its checksum dictionary for the client to compare with its local version, at which point the client can take action if it discovers a mismatch.

The dictionary type is defined in the file `Ice/SliceChecksumDict.ice` as follows:

Slice

```
module Ice {
    dictionary<string, string> SliceChecksumDict;
};
```

This type can be incorporated into an application's Slice definitions like this:

Slice

```
#include <Ice/SliceChecksumDict.ice>

interface MyServer {
    idempotent Ice::SliceChecksumDict getSliceChecksums();
    // ...
};
```

The key of each element in the dictionary is a Slice type ID, and the value is the checksum of that type.

For more information on generating and using Slice checksums, see the appropriate language mapping chapter.

See Also

- Type IDs

Copyright 2016, ZeroC, Inc.

Generating Slice Documentation

On this page:

- Generating Slice Documentation
- Documentation Comments
  - Hyperlinks
  - Explicit Cross-References
  - Markup for Operations
  - General HTML Markup
- Using slice2html

Generating Slice Documentation

The Slice API reference documents all of the Slice definitions for Ice and its services by extracting comments from their Slice files. You can generate HTML documentation from your own Slice definitions using `slice2html`. As an example of the comment syntax supported by `slice2html`, here is the definition of `Ice::Current`:

Slice

```
/**
 *
 * Information about the current method invocation for servers. Each
 * operation on the server has a <tt>Current</tt> as its implicit final
 * parameter. <tt>Current</tt> is mostly used for Ice services. Most
 * applications ignore this parameter.
 *
 **/
local struct Current {
    /**
     * The object adapter.
     **/
    ObjectAdapter adapter;

    /**
     * Information about the connection over which the current
     * method invocation was received. If the invocation is direct
     * due to collocation optimization, this value is set to null.
     **/
    Connection con;

    /**
     * The Ice object identity.
     **/
    Identity id;

    /**
     * The facet.
     ***/
    string facet;

    /**
     * The operation name.
     **/
    string operation;

    /**
```


Copyright 2016, ZeroC, Inc.

```
     * The mode of the operation.
     **/
    OperationMode mode;

    /**
     * The request context, as received from the client.
     **/
    Context ctx;

    /**
     * The request id unless oneway (0) or collocated (-1).
     **/
    int requestId;

    /**
     * The encoding version used to encode the input and output parameters.
     **/
```


Copyright 2016, ZeroC, Inc.

```
    Ice::EncodingVersion encoding;
};
```

If you look at the comments, you will see these reflected in the documentation for `Ice::Current` in the online Slice API reference.

Documentation Comments

A documentation comment:

- starts with `/**`
- ends with `**/`

Such a comment can precede any Slice construct, such as a module, interface, structure, operation, and so on. Within a documentation comment, you can either start each line with a `*`, or you can leave the beginning of the line blank — `slice2html` can handle either convention:

Slice

```
/**
 *
 * This is a documentation comment for which every line
 * starts with a '*' character.
 **/

/**

 This is a documentation comment without a leading '*'
 for each line. Either style of comment is fine.

 **/
```

The first sentence of the documentation comment for a Slice construct should be a summary sentence. `slice2html` generates an index of all Slice constructs; the first sentence of the comments for each Slice construct is ued as a summary in that index.

Hyperlinks

Any Slice identifier enclosed in `{@link ...}` is presented as a hyperlink in code font. For example:

Slice

```
/**
 * An empty {@link name} denotes a null object.
 **/
```

This generates a hyperlink for the `name` markup that points at the definition of the corresponding Slice symbol. (The symbol can denote any Slice construct, such as a type, interface, parameter, or structure member.)

Explicit Cross-References

The directive `@see` is recognized by `slice2html`. Where it appears, the generated HTML contains a separate section titled "See Also", followed by a list of Slice identifiers. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
/**
 * The object adapter, which is responsible for receiving requests
 * from endpoints, and for mapping between servants, identities,
 * and proxies.
 *
 * @see Communicator
 * @see ServantLocator
 **/
```

The Slice identifiers are listed in the corresponding "See Also" section as hyperlinks in code font.

Markup for Operations

There are three directives specifically to document Slice operations: `@param`, `@return`, and `@throws`. For example:

Slice

```
/**
 * Look for an item with the specified
 * primary and secondary key.
 *
 * @param p The primary search key.
 *
 * @param s The secondary search key.
 *
 * @return The item that matches the specified keys.
 *
 * @throws NotFound Raised if no item matches the specified keys.
 **/

Item findItem(Key p, Key s) throws NotFound;
```

`slice2html` generates separate "Parameters", "Return Value", and "Exceptions" sections for these directives. Parameters are listed in the same order as they appear in the comments. (For clarity, that order should match the order of declaration of parameters for the corresponding operation.)

General HTML Markup

A documentation comment can contain any markup that is permitted by HTML in that place. For example, you can create separate paragraphs with `<P>` and `</P>` elements:

Slice

```
/**
 * This is a comment for some Slice construct.</p>
 *
 * <p>This comment appears in a separate paragraph.
 **/
```


Copyright 2016, ZeroC, Inc.

Note that you must neither begin a documentation comment with a element nor end it with a element because, in the generated HTML, documentation comments are already surrounded by and elements.

There are various other ways to create markup — for example, you can use <table> or <ul> elements. Please see the HTML specification for details.

Using `slice2html`

`slice2html` uses the following syntax:

```
slice2html [options] slice_file...
```

If you have cross-references that span Slice files, you must compile all of the Slice files with a single invocation of `slice2html`.

The command supports the following options:

- `-h, --help`
  Displays a help message.
- `-v, --version`
  Displays the compiler version.
- `-D`*`NAME`*
  Defines the preprocessor symbol *`NAME`*.
- `-D`*`NAME=DEF`*
  Defines the preprocessor symbol *`NAME`* with the value *`DEF`*.
- `-U`*`NAME`*
  Undefines the preprocessor symbol {*`NAME`*.
- `-I`*`DIR`*
  Add the directory *`DIR`* to the search path for `#include` directives.
- `-E`
  Print the preprocessor output on `stdout`.
- `--output-dir` *`DIR`*
  Place the generated files into directory *`DIR`*.
- `-d, --debug`
  Print debug information showing the operation of the Slice parser.
- `--ice`
  Permit use of the normally reserved prefix `Ice` for identifiers. Use this option only when compiling the source code for the Ice run time.
- `--underscore`
  Permit use of underscores in Slice identifiers.
- `--hdr` *`FILE`*
  Prepend *`FILE`* to each generated HTML file (except for `_sindex.html`). This allows you to replace the HTML header and other preamble information with a custom version, so you can connect style sheets to the generated pages. The specified file must include the <body> tag (but need not end with a <body> tag). *`FILE`* is expected to contain the string `TITLE` on a line by itself, starting in column one. slice2html replaces the `TITLE` string with the fully-scoped name of the Slice symbol that is documented on the corresponding page.
- `--ftr` *`FILE`*
  Append *`FILE`* to each generated HTML file (except for `_sindex.html`). This allows you to add, for example, a custom footer to each generated page. *`FILE`* must end with a </body> tag.
- `--indexhdr` *`FILE`*
  `slice2html` generates a file `_sindex.html` that contains a table of contents of all Slice symbols that hyperlink to the corresponding page. This option allows you to replace the standard header with a custom header, for example, to attach a JavaScript. The specified file must include the <body> tag (but need not end with a <body> tag). The default value is the setting of `--hdr` (if any).

Copyright 2016, ZeroC, Inc.

- `--indexftr` *`FILE`*
  Append *`FILE`* to the generated *`sindex.html`* page. *This allows you to add, for example, a custom footer to the table of contents, or to invoke a JavaScript.* *`_FILE`* is must end with a `</body>` tag. The default value is the setting of `--ftr` (if any).

- `--image-dir` *`DIR`*
  With this option, `slice2html` looks in the specified directory for images to use for the generated navigation hyperlinks. (Without this option, text links are used instead.) Please see the generated HTML for the names of the various image files. (They can easily be found by looking for `img` elements.)

- `--logo-url` *`URL`*
  Use the specified URL as a hyperlink for the company logo that is added to each page (if `--image-dir` is specified). The company logo is expected to be in `<`*`image_dir`*`>/logo.gif`.

- `--search` *`ACTION`*
  If this option is specified, the generated pages contain a search box that allows you to connect the generated pages to a search engine. On pressing the "Search" button, the specified *`ACTION`* is carried out.

- `--index` *`NUM`*
  `slice2html` generates sub-indexes for various Slice symbols. This option controls how many entries must be present before a sub-index is generated. For example, if *`NUM`* is set to 3, a sub-index will be generated only if there are three or more symbols that appear in that index. The default settings is 1, meaning that a sub-index is always generated. To disable sub-indexes entirely, set *`NUM`* to 0.

- `--summary` *`NUM`*
  If this option is set, summary sentences that exceed *`NUM`* characters generate a warning.

See Also

- Slice API reference
- HTML specification

Copyright 2016, ZeroC, Inc.

Slice Keywords

The following identifiers are Slice keywords:

`bool`	`extends`	`LocalObject`	`string`
`byte`	`false`	`long`	`struct`
`class`	`float`	`module`	`throws`
`const`	`idempotent`	`Object`	`true`
`dictionary`	`implements`	`optional`	`void`
`double`	`int`	`out`	
`enum`	`interface`	`sequence`	
`exception`	`local`	`short`	

Keywords must be capitalized as shown.

See Also

- Lexical Rules

Copyright 2016, ZeroC, Inc.

Slice Metadata Directives

On this page:

- General Metadata Directives
- Metadata Directives for C++
- Metadata Directives for Java
- Metadata Directives for C#
- Metadata Directives for .NET
- Metadata Directives for Objective-C
- Metadata Directives for Python
- Metadata Directives for Freeze

General Metadata Directives

`amd`

This directive applies to interfaces, classes, and individual operations. It enables code generation for asynchronous method dispatch. (See the relevant language mapping chapter for details.)

`deprecate`

This directive allows you to emit a deprecation warning for Slice constructs.

`format`

This directive defines the encoding format used for any classes or exceptions marshaled as the arguments or results of an operation. The tag can be applied to an interface, which affects all of its operations, or to individual operations. Legal values for the tag are `format:sliced`, `format:compact`, and `format:default`. A tag specified for an operation overrides any setting applied to its enclosing interface. The `Ice.Default.SlicedFormat` property defines the behavior when no tag is present.

`preserve-slice`

This directive applies to classes and exceptions, allowing an intermediary to forward an instance of the annotated type, or any of its subtypes, with all of its slices intact. Operations that transfer such types must be annotated with `format:sliced`. It is not necessary to repeat the `preserve-slice` tag on derived types, but you may wish to do so for documentation purposes.

`protected`

This directive applies to data members of classes and changes code generation to make these members protected. See class mapping of the relevant language mapping chapter for more information.

`UserException`

This directive applies only to operations on local interfaces. The metadata directive indicates that the operation can throw any user exception, regardless of its specific definition. (This directive is used for the `locate` and `finished` operations on servant locators, which can throw any user exception.)

Metadata Directives for C++

`cpp:array` and `cpp:range`

These directives apply to sequences. They direct the code generator to create zero-copy APIs for passing sequences as parameters.

`cpp:class`

This directive applies to structures. It directs the code generator to create a C++ class (instead of a C++ structure) to represent a Slice structure.

`cpp:comparable`

This directive applies to structures. It directs the code generator to generate comparison operators for a structure regardless of whether it

Copyright 2016, ZeroC, Inc.

qualifies as a legal dictionary key type.

cpp:const

This directive applies to operations. It directs the code generator to create a `const` pure virtual member function for the skeleton class.

cpp:header-ext

This global directive allows you to use a file extension for C++ header files other than the default `.h` extension.

cpp:ice_print

This directive applies to exceptions. It directs the code generator to declare (but not implement) an `ice_print` member function that overrides the `ice_print` virtual function in `Ice::Exception`. The application must provide the implementation of this `ice_print` functio n.

cpp:include

This global directive allows you inject additional #include directives into the generated code. This is useful for custom types.

cpp:type:*c++-type*

This directive applies to sequences and dictionaries. It directs the code generator to map the Slice type or parameter to the provided C++ type.

cpp:type:string and cpp:type:wstring

These directives apply to data members of type string as well as to containers, such as structures, classes, exceptions, and modules. String members map by default to `std::string`. You can use the `cpp:type:wstring` metadata to cause a string data member (or all string data members in a structure, class or exception) to map to `std::wstring` instead. Use the `cpp:type:string` metadata to force string members to use the default mapping regardless of any enclosing metadata.

Slice

```
["cpp:type:wstring"]
module A { // All string members in this module map by default to std::wstring
    struct Struct1 {
        string s; // Maps to std::wstring
    };
    struct Struct2 {
        ["cpp:type:string"] string s; // Maps to std::string
    };

    ["cpp:type:string"] // All string members in this module map by default to
std::string
    module Inner {
        struct Struct4 {
            string s; // Maps to std::string
        };

        ["cpp:type:wstring"] // All string members of Struct4 map by default to
std::wstring
        struct Struct3 {
            string s; // Maps to std::wstring
        };
    };
};
```


Copyright 2016, ZeroC, Inc.

cpp:view-type:*c++-view-type*

This directive applies to string, sequence and dictionary parameters. It directs the code generator to map this parameter to the provided C++ type when this parameter does not need to hold any memory, for example when mapping an in-parameter to a proxy function.

cpp:virtual

This directive applies to classes. If the directive is present and a class has base classes, the generated C++ class derives virtually from its bases; without this directive, slice2cpp generates the class so it derives non-virtually from its bases.

This directive is useful if you use Slice classes as servants and want to inherit the implementation of operations in the base class in the derived class. For example:

Slice

```
class Base {
    int baseOp();
};

["cpp:virtual"]
class Derived extends Base {
    string derivedOp();
};
```

The metadata directive causes slice2cpp to generate the class definition for `Derived` using virtual inheritance:

C++

```
class Base : virtual public Ice::Object {
    // ...
};

class Derived : virtual public Base {
    // ...
};
```

This allows you to reuse the implementation of `baseOp` in the servant for `Derived` using ladder inheritance:

C++

```
class BaseI : public virtual Base {
    Ice::Int baseOp(const Ice::Current&);
    // ...
};

class DerivedI : public virtual Derived, public virtual BaseI {
    // Re-use inherited baseOp()
};
```

Note that, if you have data member in classes and use virtual inheritance, you need to take care to correctly call base class constructors if you implement your own one-shot constructor. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
class Base {
    int baseInt;
};

class Derived extends Base {
    int derivedInt;
};
```

The generated one-shot constructor for `Derived` initializes both `baseInt` and `derivedInt`:

C++

```
Derived::Derived(Ice::Int __ice_baseInt, Ice::Int __ice_derivedInt)
    : M::Base(__ice_baseInt),
      derivedInt(__ice_derivedInt)
{
}
```

If you derive your own class from `Derived` and add a one-shot constructor to your class, you must explicitly call the constructor of all the base classes, including `Base`. Failure to call the `Base` constructor will result in `Base` being default-constructed instead of getting a defined value. For example:

C++

```
class DerivedI : public virtual Derived {
public:
    DerivedI(int baseInt, int derivedInt, const string& s)
        : Base(baseInt), Derived(baseInt, derivedInt), _s(s)
    {
    }

private:
    string _s;
};
```

This code correctly initializes the `baseInt` member of the `Base` part of the class. Note that the following does not work as intended and leaves the `Base` part default-constructed (meaning that `baseInt` is not initialized):

Copyright 2016, ZeroC, Inc.

C++

```
class DerivedI : public virtual Derived {
public:
    DerivedI(int baseInt, int derivedInt, const string& s)
        : Derived(baseInt, derivedInt), _s(s)
    {
        // WRONG: Base::baseInt is not initialized.
    }

private:
    string _s;
};
```

Metadata Directives for Java

`java:buffer`

This directive applies to sequences of certain primitive types. It directs the translator to map the sequence to a subclass of `java.nio.Buffer`.

`java:getset`

This directive applies to data members and structures, classes, and exceptions. It adds accessor and modifier methods (JavaBean methods) for data members.

`java:optional`

This directive forces optional output parameters to use the optional mapping instead of the default required mapping in servants.

`java:package`

This global directive instructs the code generator to place the generated classes into a specific package.

`java:serializable`

This directive allows you to use Ice to transmit serializable Java classes as native objects, without having to define corresponding Slice definitions for these classes.

`java:serialVersionUID`

This directive overrides the default (generated) value of `serialVersionUID` for a Slice type.

`java:type`

This directive allows you to use custom types for sequences and dictionaries.

Metadata Directives for C#

Note that C# (and other Common Language Runtime languages) are also affected by metadata with a `clr:` prefix. (See Metadata Directives for .NET.)

`cs:attribute`

This directive can be used both as a global directive and as directive for specific Slice constructs. It injects C# attribute definitions into the generated code. (See C-Sharp Specific Metadata Directives.)

Copyright 2016, ZeroC, Inc.

Metadata Directives for .NET

`clr:class`

This directive applies to Slice structures. It directs the code generator to emit a C# class instead of a structure.

`clr:generic:List`, `clr:generic:LinkedList`, `clr:generic:Queue` and `clr:generic:Stack`

These directives apply to sequences and map them to the specified sequence type.

`clr:generic:SortedDictionary`

This directive applies to dictionaries and maps them to `SortedDictionary`.

`clr:generic`

This directive applies to sequences and allows you map them to custom types.

`clr:implements:`*`type`*

This directive adds the specified base type to the generated code for a Slice structure, class or interface. For example, Ice defines the `Communicator` interface as shown below:

Slice

```
["clr:implements:_System.IDisposable"]
local interface Communicator { ... };
```

Consequently, the generated C# interface `Ice.Communicator` implements `IDisposable`.

Every Slice-generated C# source file defines two namespace aliases:

```
using _System = global::System;
using _Microsoft = global::Microsoft;
```

We recommend using these aliases if your metadata refers to the `System` or `Microsoft` namespaces.

`clr:property`

This directive applies to Slice structures and classes. It directs the code generator to create C# property definitions for data members.

`clr:serializable`

This directive allows you to use Ice to transmit serializable CLR classes as native objects, without having to define corresponding Slice definitions for these classes.

Metadata Directives for Objective-C

`objc:prefix`

This directive applies to modules and changes the default mapping for modules to use a specified prefix.

Metadata Directives for Python

`python:package`

Copyright 2016, ZeroC, Inc.

This global directive instructs the code generator to place the generated code into a specified Python package.

`python:seq:default`, `python:seq:list` and `python:seq:tuple`

These directives allow you to change the mapping for Slice sequences.

Metadata Directives for Freeze

`freeze:read` and `freeze:write`

These directives inform a Freeze evictor whether an operation updates the state of an object, so the evictor knows whether it must save an object before evicting it.

See Also

- Metadata

Copyright 2016, ZeroC, Inc.

Slice for a Simple File System

For this manual, we use a file system application to illustrate various aspects of Ice. Throughout, we progressively improve and modify the application such that it evolves into an application that is realistic and illustrates the architectural and coding aspects of Ice. This allows us to explore the capabilities of the platform to a realistic degree of complexity without overwhelming you with an inordinate amount of detail early on.

In this section:

- File System Application outlines the file system functionality
- Slice Definitions for the File System develops the data types and interfaces that are required for the file system
- Complete Definition presents the complete Slice definition for the application.

File System Application

Our file system application implements a simple hierarchical file system, similar to the file systems we find in Windows or Unix. To keep code examples to manageable size, we ignore many aspects of a real file system, such as ownership, permissions, symbolic links, and a number of other features. However, we build enough functionality to illustrate how you could implement a fully-featured file system, and we pay attention to things such as performance and scalability. In this way, we can create an application that presents us with real-world complexity without getting buried in large amounts of code.

Our file system consists of directories and files. Directories are containers that can contain either directories or files, meaning that the file system is hierarchical. A dedicated directory is at the root of the file system. Each directory and file has a name. Files and directories with a common parent directory must have different names (but files and directories with different parent directories can have the same name). In other words, directories form a naming scope, and entries with a single directory must have unique names. Directories allow you to list their contents.

For now, we do not have a concept of pathnames, or the creation and destruction of files and directories. Instead, the server provides a fixed number of directories and files. (We will address the creation and destruction of files and directories in Object Life Cycle.)

Files can be read and written but, for now, reading and writing always replace the entire contents of a file; it is impossible to read or write only parts of a file.

Slice Definitions for the File System

Given the very simple requirements we just outlined, we can start designing interfaces for the system. Files and directories have something in common: they have a name and both files and directories can be contained in directories. This suggests a design that uses a base type that provides the common functionality, and derived types that provide the functionality specific to directories and files, as shown:

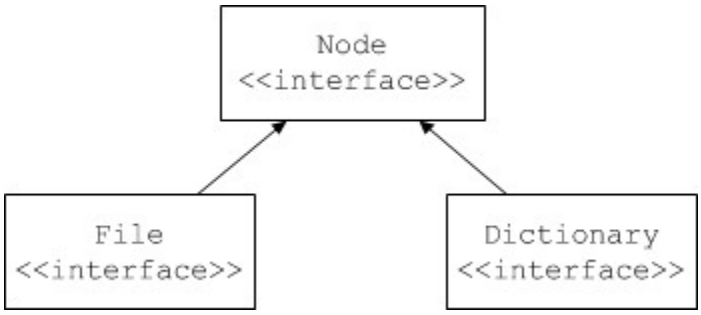

Inheritance Diagram of the File System.

The Slice definitions for this look as follows:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Node {
    // ...
};

interface File extends Node {
    // ...
};

interface Directory extends Node {
    // ...
};
```

Next, we need to think about what operations should be provided by each interface. Seeing that directories and files have names, we can add an operation to obtain the name of a directory or file to the `Node` base interface:

Slice

```
interface Node {
    idempotent string name();
};
```

The `File` interface provides operations to read and write a file. For simplicity, we limit ourselves to text files and we assume that `read` operations never fail and that only `write` operations can encounter error conditions. This leads to the following definitions:

Slice

```
exception GenericError {
    string reason;
};

sequence<string> Lines;

interface File extends Node {
    idempotent Lines read();
    idempotent void write (Lines text) throws GenericError;
};
```

Note that `read` and `write` are marked idempotent because either operation can safely be invoked with the same parameter value twice in a row: the net result of doing so is the same has having (successfully) called the operation only once.

The `write` operation can raise an exception of type `GenericError`. The exception contains a single `reason` data member, of type `string`. If a `write` operation fails for some reason (such as running out of file system space), the operation throws a `GenericError` exception, with an explanation of the cause of the failure provided in the `reason` data member.

Directories provide an operation to list their contents. Because directories can contain both directories and files, we take advantage of the polymorphism provided by the `Node` base interface:

Copyright 2016, ZeroC, Inc.

Slice

```
sequence<Node*> NodeSeq;

interface Directory extends Node {
    idempotent NodeSeq list();
};
```

The `NodeSeq` sequence contains elements of type `Node*`. Because `Node` is a base interface of both `Directory` and `File`, the `NodeSeq` sequence can contain proxies of either type. (Obviously, the receiver of a `NodeSeq` must down-cast each element to either `File` or `Directory` in order to get at the operations provided by the derived interfaces; only the `name` operation in the `Node` base interface can be invoked directly, without doing a down-cast first. Note that, because the elements of `NodeSeq` are of type `Node*` (not `Node`), we are using pass-by-reference semantics: the values returned by the `list` operation are proxies that each point to a remote node on the server.

These definitions are sufficient to build a simple (but functional) file system. Obviously, there are still some unanswered questions, such as how a client obtains the proxy for the root directory. We will address these questions in the relevant implementation chapter.

Complete Definition

We wrap our definitions in a module, resulting in the final definition as follows:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };

    exception GenericError {
        string reason;
    };

    sequence<string> Lines;

    interface File extends Node {
        idempotent Lines read();
        idempotent void write(Lines text) throws GenericError;
    };

    sequence<Node*> NodeSeq;

    interface Directory extends Node {
        idempotent NodeSeq list();
    };
};
```

See Also

- Object Life Cycle

Copyright 2016, ZeroC, Inc.

Language Mappings

Topics

- C++ Mapping
- C-Sharp Mapping
- Java Mapping
- JavaScript Mapping
- Objective-C Mapping
- PHP Mapping
- Python Mapping
- Ruby Mapping

Copyright 2016, ZeroC, Inc.

C++ Mapping

Topics

- Client-Side Slice-to-C++ Mapping
- Server-Side Slice-to-C++ Mapping
- Customizing the C++ Mapping
- Version Information in C++
- slice2cpp Command-Line Options
- C++ Strings and Character Encoding
- The C++ Utility Library

Copyright 2016, ZeroC, Inc.

Client-Side Slice-to-C++ Mapping

The client-side Slice-to-C++ mapping defines how Slice data types are translated to C++ types, and how clients invoke operations, pass parameters, and handle errors. Much of the C++ mapping is intuitive. For example, Slice sequences map to STL vectors, so there is essentially nothing new you have to learn in order to use Slice sequences in C++.

The rules that make up the C++ mapping are simple and regular. In particular, the mapping is free from the potential pitfalls of memory management: all types are self-managed and automatically clean up when instances go out of scope. This means that you cannot accidentally introduce a memory leak by, for example, ignoring the return value of an operation invocation or forgetting to deallocate memory that was allocated by a called operation.

The C++ mapping is fully thread-safe. For example, the reference counting mechanism for classes is interlocked against parallel access, so reference counts cannot be corrupted if a class instance is shared among a number of threads. Obviously, you must still synchronize access to data from different threads. For example, if you have two threads sharing a sequence, you cannot safely have one thread insert into the sequence while another thread is iterating over the sequence. However, you only need to concern yourself with concurrent access to your own data — the Ice run time itself is fully thread safe, and none of the Ice API calls require you to acquire or release a lock before you safely can make the call.

Much of what appears in this chapter is reference material. We suggest that you skim the material on the initial reading and refer back to specific sections as needed. However, we recommend that you read at least the mappings for exceptions, interfaces, and operations in detail because these sections cover how to call operations from a client, pass parameters, and handle exceptions.

In order to use the C++ mapping, you should need no more than the Slice definition of your application and knowledge of the C++ mapping rules. In particular, looking through the generated header files in order to discern how to use the C++ mapping is likely to be confusing because the header files are not necessarily meant for human consumption and, occasionally, contain various cryptic constructs to deal with operating system and compiler idiosyncrasies. Of course, occasionally, you may want to refer to a header file to confirm a detail of the mapping, but we recommend that you otherwise use the material presented here to see how to write your client-side code.

The `Ice` Namespace
All of the APIs for the Ice run time are nested in the `Ice` namespace, to avoid clashes with definitions for other libraries or applications. Some of the contents of the `Ice` namespace are generated from Slice definitions; other parts of the `Ice` namespace provide special-purpose definitions that do not have a corresponding Slice definition. We will incrementally cover the contents of the `Ice` namespace throughout the remainder of the manual.

Topics

- C++ Mapping for Identifiers
- C++ Mapping for Modules
- C++ Mapping for Built-In Types
- C++ Mapping for Enumerations
- C++ Mapping for Structures
- C++ Mapping for Sequences
- C++ Mapping for Dictionaries
- C++ Mapping for Constants
- C++ Mapping for Exceptions
- C++ Mapping for Interfaces
- C++ Mapping for Operations
- C++ Mapping for Optional Values
- C++ Mapping for Classes
- Smart Pointers for Classes
- Asynchronous Method Invocation (AMI) in C++
- Using Slice Checksums in C++
- Example of a File System Client in C++

Copyright 2016, ZeroC, Inc.

C++ Mapping for Identifiers

A Slice identifier maps to an identical C++ identifier. For example, the Slice identifier `Clock` becomes the C++ identifier `Clock`. There is one exception to this rule: if a Slice identifier is the same as a C++ keyword, the corresponding C++ identifier is prefixed with `_cpp_`. For example, the Slice identifier `while` is mapped as `_cpp_while`.

A single Slice identifier often results in several C++ identifiers. For example, for a Slice interface named `Foo`, the generated C++ code uses the identifiers `Foo` and `FooPrx` (among others). If the interface has the name `while`, the generated identifiers are `_cpp_while` and `while Prx` (*not* `_cpp_whilePrx`), that is, the prefix is applied only to those generated identifiers that actually require it.

You should try to avoid such identifiers as much as possible.

See Also

- Lexical Rules
- C++ Mapping for Modules
- C++ Mapping for Built-In Types
- C++ Mapping for Enumerations
- C++ Mapping for Structures
- C++ Mapping for Sequences
- C++ Mapping for Dictionaries
- C++ Mapping for Constants
- C++ Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

C++ Mapping for Modules

A Slice module maps to a C++ namespace. The mapping preserves the nesting of the Slice definitions. For example:

Slice

```
module M1 {
    module M2 {
        // ...
    };
    // ...
};

// ...

module M1 {     // Reopen M1
    // ...
};
```

This definition maps to the corresponding C++ definition:

C++

```
namespace M1 {
    namespace M2 {
        // ...
    }
    // ...
}

// ...

namespace M1 {  // Reopen M1
    // ...
}
```

If a Slice module is reopened, the corresponding C++ namespace is reopened as well.
See Also

- Modules
- C++ Mapping for Identifiers
- C++ Mapping for Built-In Types
- C++ Mapping for Enumerations
- C++ Mapping for Structures
- C++ Mapping for Sequences
- C++ Mapping for Dictionaries
- C++ Mapping for Constants
- C++ Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

C++ Mapping for Built-In Types

On this page:

- Mapping of Slice Built-In Types to C++ Types
- Alternative String Mapping for C++
- String View Mapping in C++

Mapping of Slice Built-In Types to C++ Types

The Slice built-in types are mapped to C++ types as shown in this table:

Slice	C++
bool	`bool`
byte	`Ice::Byte`
short	`Ice::Short`
int	`Ice::Int`
long	`Ice::Long`
float	`Ice::Float`
double	`Ice::Double`
string	`std::string`

Slice `bool` and `string` map to C++ `bool` and `std::string`. The remaining built-in Slice types map to C++ type definitions instead of C++ native types. This allows the Ice run time to provide a definition as appropriate for each target architecture. (For example, `Ice::Int` might be defined as `long` on one architecture and as `int` on another.)

Note that `Ice::Byte` is a typedef for `unsigned char`. This guarantees that byte values are always in the range 0..255.

All the basic types are guaranteed to be distinct C++ types, that is, you can safely overload functions that differ in only the types listed in the table above.

Alternative String Mapping for C++

You can use a metadata directive, `"cpp:type:wstring"`, to map strings to C++ `std::wstring`. This is useful for applications that use languages with alphabets that cannot be represented in 8-bit characters. The metadata directive can be applied to any Slice construct. For containers (such as modules, interfaces, or structures), the metadata directive applies to all strings within the container. A corresponding metadata directive, `"cpp:type:string"`, can be used to selectively override the mapping defined by the enclosing container. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
["cpp:type:wstring"]
struct S1 {
    string x;             // Maps to std::wstring
    ["cpp:type:wstring"]
    string y;             // Maps to std::wstring
    ["cpp:type:string"]
    string z;             // Maps to std::string
};

struct S2 {
    string x;             // Maps to std::string
    ["cpp:type:string"]
    string y;             // Maps to std::string
    ["cpp:type:wstring"]
    string z;             // Maps to std::wstring
};
```

With these metadata directives, the strings are mapped as indicated by the comments. By default, narrow strings are encoded as UTF-8, and wide strings use Unicode in an encoding that is appropriate for the platform on which the application executes. You can override the encoding for narrow and wide strings by registering a string converter with the Ice run time.

String View Mapping in C++

You can use the metadata directive `cpp:view-type:`*`string-view-type`* to map some string parameters to a custom C++ "view-type" of your choice. This view-type can reference memory without owning it, like the experimental `string_view` type. For example:

Slice

```
void sendString(["cpp:view-type:std::experimental::string_view"] string data);
```

maps to:

C++

```
// Proxy function for synchronous call: input parameter mapped to string_view type.
void sendString(const std::experimental::string_view&);
```

See Customizing the C++ Mapping for a detailed description of the `cpp:view-type` metadata directive.

See Also

- Basic Types
- Customizing the C++ Mapping
- C++ Mapping for Identifiers
- C++ Mapping for Modules
- C++ Mapping for Enumerations
- C++ Mapping for Structures

Copyright 2016, ZeroC, Inc.

- C++ Mapping for Sequences
- C++ Mapping for Dictionaries
- C++ Mapping for Constants
- C++ Mapping for Exceptions
- C++ Strings and Character Encoding

Copyright 2016, ZeroC, Inc.

C++ Mapping for Enumerations

A Slice enumeration maps to the corresponding enumeration in C++. For example:

Slice

```
enum Fruit { Apple, Pear, Orange };
```

Not surprisingly, the generated C++ definition is identical:

C++

```
enum Fruit { Apple, Pear, Orange };
```

Suppose we modify the Slice definition to include a custom enumerator value:

Slice

```
enum Fruit { Apple, Pear = 3, Orange };
```

The generated C++ definition now includes an explicit initializer for every enumerator:

C++

```
enum Fruit { Apple = 0, Pear = 3, Orange = 4 };
```

See Also

- Enumerations
- C++ Mapping for Structures
- C++ Mapping for Sequences
- C++ Mapping for Dictionaries

Copyright 2016, ZeroC, Inc.

C++ Mapping for Structures

A Slice structure maps to a C++ structure by default. In addition, you can use a metadata directive to map structures to C++ classes.

On this page:

- Default Mapping for Structures in C++
- Class Mapping for Structures in C++
- Default Constructors for Structures in C++

Default Mapping for Structures in C++

Slice structures map to C++ structures with the same name. For each Slice data member, the C++ structure contains a public data member. For example, here is our Employee structure once more:

Slice

```
struct Employee {
    long number;
    string firstName;
    string lastName;
};
```

The Slice-to-C++ compiler generates the following definition for this structure:

C++

```
struct Employee {
    Ice::Long   number;
    std::string firstName;
    std::string lastName;
    bool operator==(const Employee&) const;
    bool operator!=(const Employee&) const;
    bool operator<(const Employee&) const;
    bool operator<=(const Employee&) const;
    bool operator>(const Employee&) const;
    bool operator>=(const Employee&) const;
};
```

For each data member in the Slice definition, the C++ structure contains a corresponding public data member of the same name. Constructors are intentionally omitted so that the C++ structure qualifies as a *plain old datatype* (POD).

The structure may also contain comparison operators to allow its use as the key type of Slice dictionaries, which are mapped to `std::map` in C++. These operators have the following behavior:

- `operator==`
  Two structures are equal if (recursively), all its members are equal.
- `operator!=`
  Two structures are not equal if (recursively), one or more of its members are not equal.
- `operator<`
  `operator<=`
  `operator>`
  `operator>=`
  The comparison operators treat the members of a structure as sort order criteria: the first member is considered the first criterion,

Copyright 2016, ZeroC, Inc.

the second member the second criterion, and so on. Assuming that we have two `Employee` structures, `s1` and `s2`, this means that the generated code uses the following algorithm to compare `s1` and `s2`:

C++

```
bool Employee::operator<(const Employee& rhs) const
{
    if (this == &rhs)   // Short?cut self?comparison
        return false;

    // Compare first members
    //
    if (number < rhs.number)
        return true;
    else if (rhs.number < number)
        return false;

    // First members are equal, compare second members
    //
    if (firstName < rhs.firstName)
        return true;
    else if (rhs.firstName < firstName)
        return false;

    // Second members are equal, compare third members
    //
    if (lastName < rhs.lastName)
        return true;
    else if (rhs.lastName < lastName)
        return false;

    // All members are equal, so return false
    return false;
}
```

As of Ice 3.5, the comparison operators are only generated for structures that qualify as legal dictionary keys. You can force the Slice compiler to generate the comparison operators for non-qualifying structures by using the `cpp:comparable` metadata:

Slice

```
["cpp:comparable"]
struct NoKey
{
    float f;
};
```

Note that the ordering behavior for non-qualifying structures may not be deterministic.

Copy construction and assignment always have deep-copy semantics. You can freely assign structures or

Copyright 2016, ZeroC, Inc.

structure members to each other without having to worry about memory management. The following code fragment illustrates both comparison and deep-copy semantics:

C++

```
Employee e1, e2;
e1.firstName = "Bjarne";
e1.lastName = "Stroustrup";
e2 = e1;                        // Deep copy
assert(e1 == e2);
e2.firstName = "Andrew";        // Deep copy
e2.lastName = "Koenig";         // Deep copy
assert(e2 < e1);
```

Because strings are mapped to `std::string`, there are no memory management issues in this code and structure assignment and copying work as expected. (The default member-wise copy constructor and assignment operator generated by the C++ compiler do the right thing.)

Class Mapping for Structures in C++

Occasionally, the mapping of Slice structures to C++ structures can be inefficient. For example, you may need to pass structures around in your application, but want to avoid having to make expensive copies of the structures. (This overhead becomes noticeable for structures with many complex data members, such as sequences or strings.) Of course, you could pass the structures by const reference, but that can create its own share of problems, such as tracking the life time of the structures to avoid ending up with dangling references.

For this reason, you can enable an alternate mapping that maps Slice structures to C++ classes. Classes (as opposed to structures) are reference-counted. Because the Ice C++ mapping provides smart pointers for classes, you can keep references to a class instance in many places in the code without having to worry about either expensive copying or life time issues.

The alternate mapping is enabled by a metadata directive, `["cpp:class"]`. Here is our Employee structure once again, but this time with the additional metadata directive:

Slice

```
["cpp:class"] struct Employee {
    long number;
    string firstName;
    string lastName;
};
```

Here is the generated class:

Copyright 2016, ZeroC, Inc.

C++

```
class Employee : public IceUtil::Shared {
public:
    Employee() {}
    Employee(::Ice::Long,
             const ::std::string&,
             const ::std::string&);
    ::Ice::Long number;
    ::std::string firstName;
    ::std::string lastName;

    bool operator==(const Employee&) const;
    bool operator!=(const Employee&) const;
    bool operator<(const Employee&) const;
    bool operator<=(const Employee&) const;
    bool operator>(const Employee&) const;
    bool operator>=(const Employee&) const;
};
```

Note that the generated class, apart from a default constructor, has a constructor that accepts one argument for each member of the structure. This allows you to instantiate and initialize the class in a single statement (instead of having to first instantiate the class and then assign to its members).

As for the default structure mapping, the class contains one public data member for each data member of the corresponding Slice structure.

The comparison operators behave as for the default structure mapping.

You can learn how to instantiate classes, and how to access them via smart pointers, in the sections describing the mapping for Slice classes — the API described there applies equally to Slice structures that are mapped to classes.

Default Constructors for Structures in C++

Structures have an implicit default constructor that default-constructs each data member. Members having a complex type, such as strings, sequences, and dictionaries, are initialized by their own default constructor. However, the default constructor performs no initialization for members having one of the simple built?in types boolean, integer, floating point, or enumeration. For such a member, it is not safe to assume that the member has a reasonable default value. This is especially true for enumerated types as the member's default value may be outside the legal range for the enumeration, in which case an exception will occur during marshaling unless the member is explicitly set to a legal value.

To ensure that data members of primitive types are initialized to reasonable values, you can declare default values in your Slice definition. The default constructor initializes each of these data members to its declared value.

If you declare a default value for at least one member of a structure, or use the class mapping for the structure, the Slice compiler also generates a second constructor. This *one-shot* constructor has one parameter for each data member, allowing you to construct and initialize an instance in a single statement (instead of first having to construct the instance and then assign to its members).

See Also

- Structures
- C++ Mapping for Enumerations
- C++ Mapping for Sequences
- C++ Mapping for Dictionaries

Copyright 2016, ZeroC, Inc.

C++ Mapping for Sequences

On this page:

- Default Sequence Mapping in C++
- Custom Sequence Mapping in C++
- Custom Mapping for Sequence Parameters in C++
  - Array Mapping for Sequence Parameters in C++
  - Range Mapping for Sequence Parameters in C++

Default Sequence Mapping in C++

Here is the definition of our `FruitPlatter` sequence once more:

Slice

```
sequence<Fruit> FruitPlatter;
```

The Slice compiler generates the following definition for the `FruitPlatter` sequence:

C++

```
typedef std::vector<Fruit> FruitPlatter;
```

As you can see, the sequence simply maps to a standard `std::vector`, so you can use the sequence like any other vector. For example:

C++

```
// Make a small platter with one Apple and one Orange
//
FruitPlatter p;
p.push_back(Apple);
p.push_back(Orange);
```

Custom Sequence Mapping in C++

You can override the default mapping of Slice sequences to C++ vectors with the cpp:type and cpp:view-type metadata directives.

For example:

Copyright 2016, ZeroC, Inc.

Slice

```
[["cpp:include:list"]]

module Food {

    enum Fruit { Apple, Pear, Orange };

    ["cpp:type:std::list< ::Food::Fruit>"]
    sequence<Fruit> FruitPlatter;

};
```

With this metadata directive, the sequence now maps to a C++ `std::list`:

C++

```
#include <list>

namespace Food {

    typedef std::list<Food::Fruit> FruitPlatter;

    // ...
}
```

Custom Mapping for Sequence Parameters in C++

In addition to the default and custom mappings of sequence types as a whole, you can use metadata to customize the mapping of a single operation parameter of type sequence.

Ice provides two metadata directives for this purpose, `["cpp:array"]` and `["cpp:range"]`.

Array Mapping for Sequence Parameters in C++

The array mapping for sequence parameters applies only to:

- In parameters, on the client-side and on the server-side
- Out and return parameters provided by the Ice run-time to AMI type-safe callbacks and AMI lambdas
- Out and return parameters provided to AMD callbacks

For example:

Slice

```
interface File {
    void write(["cpp:array"] Ice::ByteSeq contents);
};
```

The `cpp:array` metadata directive instructs the compiler to map the `contents` parameter to a pair of pointers. With this directive, the `write` method on the proxy has the following signature:

Copyright 2016, ZeroC, Inc.

C++

```
void write(const std::pair<const Ice::Byte*, const Ice::Byte*>& contents);
```

To pass a byte sequence to the server, you pass a pair of pointers; the first pointer points at the beginning of the sequence, and the second pointer points one element past the end of the sequence.

Similarly, for the server side, the `write` method on the skeleton has the following signature:

C++

```
virtual void write(const ::std::pair<const ::Ice::Byte*, const ::Ice::Byte*>&,
                   const ::Ice::Current& = ::Ice::Current()) = 0;
```

The passed pointers denote the beginning and end of the sequence as a range `[first, last)` (that is, they use the usual semantics for iterators).

The array mapping is useful to achieve zero-copy passing of sequences. The pointers point directly into the server-side transport buffer; this allows the server-side run time to avoid creating a `vector` to pass to the operation implementation, thereby avoiding both allocating memory for the sequence and copying its contents into that memory.

You can use the array mapping for any sequence type. However, it provides a performance advantage only for byte sequences (on all platforms) and for sequences of integral or floating point types (on some platforms).

Range Mapping for Sequence Parameters in C++

The range mapping for sequences is similar to the array mapping and exists for the same purpose, namely, to enable zero-copy of sequence parameters:

Slice

```
interface File {
    void write(["cpp:range"] Ice::ByteSeq contents);
};
```

The `cpp:range` metadata directive instructs the compiler to map the `contents` parameter to a pair of `const_iterator`. With this directive, the `write` method on the proxy has the following signature:

C++

```
void write(const std::pair<Ice::ByteSeq::const_iterator,
Ice::ByteSeq::const_iterator>& contents);
```

Similarly, for the server side, the `write` method on the skeleton has the following signature:

Copyright 2016, ZeroC, Inc.

C++

```
virtual void write(
    const ::std::pair<::Ice::ByteSeq::const_iterator,
::Ice::ByteSeq::const_iterator>&,
    const ::Ice::Current& = ::Ice::Current()) = 0;
```

The passed iterators denote the beginning and end of the sequence as a range `[first, last)` (that is, they use the usual semantics for iterators).

The motivation for the range mapping is the same as for the array mapping: the passed iterators point directly into the server-side transport buffer and so avoid the need to create a temporary `vector` to pass to the operation.

As for the array mapping, the range mapping can be used with any sequence type, but offers a performance advantage only for byte sequences (on all platforms) and for sequences of integral type (x86 platforms only).

You can optionally add a type name to the `cpp:range` metadata directive, for example:

Slice

```
interface File {
    void write(["cpp:range:std::deque<Ice::Byte>"] Ice::ByteSeq contents);
};
```

This instructs the compiler to generate a pair of `const_iterator` for the specified type:

C++

```
virtual void write(
    const ::std::pair<std::deque<Ice::Byte>::const_iterator,
                      std::deque<Ice::Byte>::const_iterator>&,
    const ::Ice::Current& = ::Ice::Current()) = 0;
```

This is useful if you want to combine the range mapping with a custom sequence type that behaves like an standard container.

See Also

- Sequences
- Customizing the C++ Mapping
- C++ Mapping for Enumerations
- C++ Mapping for Structures
- C++ Mapping for Dictionaries
- C++ Mapping for Operations

Copyright 2016, ZeroC, Inc.

C++ Mapping for Dictionaries

On this page:

- Default Dictionary Mapping in C++
- Custom Dictionary Mapping in C++

Default Dictionary Mapping in C++

Here is the definition of our `EmployeeMap` once more:

Slice

```
dictionary<long, Employee> EmployeeMap;
```

The following code is generated for this definition:

C++

```
typedef std::map<Ice::Long, Employee> EmployeeMap;
```

Again, there are no surprises here: a Slice dictionary simply maps to a standard `std::map`. As a result, you can use the dictionary like any other `map`, for example:

C++

```
EmployeeMap em;
Employee e;

e.number = 42;
e.firstName = "Stan";
e.lastName = "Lippman";
em[e.number] = e;

e.number = 77;
e.firstName = "Herb";
e.lastName = "Sutter";
em[e.number] = e;
```

Custom Dictionary Mapping in C++

You can override the default mapping of Slice dictionaries to C++ maps with a `cpp:type` a or `cpp:view-type` metadata directive, for example:

Copyright 2016, ZeroC, Inc.

Slice

```
[["cpp:include:unordered_map"]]

["cpp:type:std::unordered_map<Ice::Long, Employee>"] dictionary<long, Employee>
EmployeeMap;
```

With this metadata directive, the dictionary now maps to a C++ `std::unordered_map`:

C++

```
#include <unordered_map>

typedef std::unordered_map<Ice::Long, Employee> EmployeeMap;
```

See Customizing the C++ Mapping for detailed information about the `cpp:type` and `cpp:view-type` metadata directives.

See Also

- Dictionaries
- Customizing the C++ Mapping
- C++ Mapping for Enumerations
- C++ Mapping for Structures
- C++ Mapping for Sequences

Copyright 2016, ZeroC, Inc.

C++ Mapping for Constants

Slice constant definitions map to corresponding C++ constant definitions. For example:

Slice

```
const bool      AppendByDefault = true;
const byte      LowerNibble = 0x0f;
const string    Advice = "Don't Panic!";
const short     TheAnswer = 42;
const double    PI = 3.1416;

enum Fruit { Apple, Pear, Orange };
const Fruit     FavoriteFruit = Pear;
```

Here are the generated definitions for these constants:

C++

```
const bool          AppendByDefault = true;
const Ice::Byte     LowerNibble =     15;
const std::string   Advice =          "Don't Panic!";
const Ice::Short    TheAnswer =       42;
const Ice::Double   PI =              3.1416;

enum Fruit { Apple, Pear, Orange };
const Fruit         FavoriteFruit =   Pear;
```

All constants are initialized directly in the header file, so they are compile-time constants and can be used in contexts where a compile-time constant expression is required, such as to dimension an array or as the `case` label of a `switch` statement.

A Slice string literal that contains non-ASCII characters is mapped by default to a narrow C++ string literal with the non-ASCII characters replaced by their UTF-8 encoding in octal escapes. For example:

Slice

```
const string Egg = "œuf";
```

is mapped to:

C++

```
const std::string Egg = "\305\223uf";
```

If you map a string constant to a `std::wstring`, the non-ASCII characters in the string literal are replaced by universal character names, except for non-ASCII characters in the range `\u0000` to `\u009f` that are replaced by octal escapes with the same value. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
const ["cpp:type:wstring"] string LargeEgg = "gros œuf";
```

is mapped to:

C++

```
const std::wstring LargeEgg = L"gros \u0153uf";
```

Likewise, a Slice string literal that contains universal character names is mapped to a narrow C++ string with the universal character names r eplaced by their UTF-8 encoding in octal escapes, or to a wide C++ string with the universal character names preserved (except for characters in the range `\u0000` to `\u009f` that are replaced by octal escapes with the same value). For example:

Slice

```
const string Heart = "c\u0153ur";
const ["cpp:type:wstring"] string BigHeart = "grand c\u0153ur";
const ["cpp:type:wstring"] string Banana = "\U0001F34C";

const string DoubleTilde = "~\u007e";
const ["cpp:type:wstring"] string WDoubleTilde = "~\u007e";
```

is mapped to:

C++

```
const std::string Heart = "c\305\223ur";
const std::wstring BigHeart = L"grand c\u0153ur";
const std::wstring Banana = L"\U0001F34C";

const std::string DoubleTilde = "~\176";
const std::string WDoubleTilde = L"~\176";
```

See Also

- Constants and Literals
- C++ Mapping for Identifiers
- C++ Mapping for Modules
- C++ Mapping for Built-In Types
- C++ Mapping for Enumerations
- C++ Mapping for Structures
- C++ Mapping for Sequences
- C++ Mapping for Dictionaries
- C++ Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

C++ Mapping for Exceptions

On this page:

- C++ Mapping for User Exceptions
- C++ Default Constructors for Exceptions
- C++ Mapping for Run-Time Exceptions

C++ Mapping for User Exceptions

Here is a fragment of the Slice definition for our world time server once more:

Slice

```
exception GenericError {
    string reason;
};
exception BadTimeVal extends GenericError {};
exception BadZoneName extends GenericError {};
```

These exception definitions map as follows:

Copyright 2016, ZeroC, Inc.

C++

```
class GenericError: public Ice::UserException {
public:
    std::string reason;

    GenericError() {}
    explicit GenericError(const string&);

    virtual const std::string&  ice_name() const;
    virtual Ice::Exception*     ice_clone() const;
    virtual void                ice_throw() const;
    // Other member functions here...
};

class BadTimeVal: public GenericError {
public:
    BadTimeVal() {}
    explicit BadTimeVal(const string&);

    virtual const std::string&  ice_name() const;
    virtual Ice::Exception*     ice_clone() const;
    virtual void                ice_throw() const;
    // Other member functions here...
};

class BadZoneName: public GenericError {
public:
    BadZoneName() {}
    explicit BadZoneName(const string&);

    virtual const std::string&  ice_name() const;
    virtual Ice::Exception*     ice_clone() const;
    virtual void                ice_throw() const;
};
```

Each Slice exception is mapped to a C++ class with the same name. For each exception member, the corresponding class contains a public data member. (Since `BadTimeVal` and `BadZoneName` do not have members, the generated classes for these exceptions also do not have members.) Optional data members are mapped to instances of the `IceUtil::Optional` template.

The inheritance structure of the Slice exceptions is preserved for the generated classes, so `BadTimeVal` and `BadZoneName` inherit from `GenericError`.

Each exception has three additional member functions:

- `ice_name`
  As the name suggests, this member function returns the name of the exception. For example, if you call the `ice_name` member function of a `BadZoneName` exception, it (not surprisingly) returns the string `"BadZoneName"`. The `ice_name` member function is useful if you catch exceptions generically and want to produce a more meaningful diagnostic, for example:

Copyright 2016, ZeroC, Inc.

C++

```
try {
    // ...
} catch(const GenericError& e) {
    cerr << "Caught an exception: " << e.ice_name() << endl;
}
```

If an exception is raised, this code prints the name of the actual exception (`BadTimeVal` or `BadZoneName`) because the exception is being caught by reference (to avoid slicing).

- `ice_clone`
  This member function allows you to polymorphically clone an exception. For example:

  C++

  ```
  try {
      // ...
  } catch(const Ice::UserException& e) {
     Ice::UserException* copy = e.clone();
  }
  ```

  `ice_clone` is useful if you need to make a copy of an exception without knowing its precise run-time type. This allows you to remember the exception and throw it later by calling `ice_throw`.

- `ice_throw`
  `ice_throw` allows you to throw an exception without knowing its precise run-time type. It is implemented as:

  C++

  ```
  void
  GenericError::ice_throw() const
  {
      throw *this;
  }
  ```

  You can call `ice_throw` to throw an exception that you previously cloned with `ice_clone`.

Each exception has a default constructor. Members having a complex type, such as strings, sequences, and dictionaries, are initialized by their own default constructor. However, the default constructor performs no initialization for members having one of the simple built-in types boolean, integer, floating point, or enumeration. For such a member, it is not safe to assume that the member has a reasonable default value. This is especially true for enumerated types as the member's default value may be outside the legal range for the enumeration, in which case an exception will occur during marshaling unless the member is explicitly set to a legal value.

To ensure that data members of primitive types are initialized to reasonable values, you can declare default values in your Slice definition. The default constructor initializes each of these data members to its declared value. Optional data members are unset unless they declare default values.

An exception also has a second constructor that accepts one argument for each exception member. This constructor allows you to instantiate and initialize an exception in a single statement, instead of having to first instantiate the exception and then assign to its members. For each optional data member, its corresponding constructor parameter uses the same mapping as for operation parameters, allowing you to pass its initial value or `IceUtil::None` to indicate an unset value.

For derived exceptions, the constructor accepts one argument for each base exception member, plus one argument for each derived exception member, in base-to-derived order.

Copyright 2016, ZeroC, Inc.

Note that the generated exception classes contain other member functions that are not shown here. However, those member functions are internal to the C++ mapping and are not meant to be called by application code.

All user exceptions ultimately inherit from `Ice::UserException`. In turn, `Ice::UserException` inherits from `Ice::Exception` (which is an alias for `IceUtil::Exception`):

C++

```
namespace IceUtil {
    class Exception : public std::exception {
        virtual const std::string& ice_name() const;
        Exception*                 ice_clone() const;
        void                       ice_throw() const;
        virtual void               ice_print(std::ostream&) const;
        // ...
    };
    std::ostream& operator<<(std::ostream&, const Exception&);
    // ...
}

namespace Ice {
    typedef IceUtil::Exception Exception;

    class UserException: public Exception {
    public:
        virtual const std::string& ice_name() const = 0;
        // ...
    };
}
```

`Ice::Exception` forms the root of the exception inheritance tree. Apart from the usual `ice_name`, `ice_clone`, and `ice_throw` member functions, it contains the `ice_print` member functions. `ice_print` prints the name of the exception. For example, calling `ice_print` on a `BadTimeVal` exception prints:

```
BadTimeVal
```

To make printing more convenient, `operator<<` is overloaded for `Ice::Exception`, so you can also write:

C++

```
try {
    // ...
} catch(const Ice::Exception& e) {
    cerr << e << endl;
}
```

This produces the same output because `operator<<` calls `ice_print` internally. You can optionally provide your own `ice_print` implem entation using the `cpp::ice_print` metadata directive.

For Ice run time exceptions, `ice_print` also shows the file name and line number at which the exception was thrown.

C++ Default Constructors for Exceptions

Exceptions have a default constructor that default-constructs each data member. Members having a complex type, such as strings, sequences, and dictionaries, are initialized by their own default constructor. However, the default constructor performs no initialization for

Copyright 2016, ZeroC, Inc.

members having one of the simple built-in types boolean, integer, floating point, or enumeration. For such a member, it is not safe to assume that the member has a reasonable default value. This is especially true for enumerated types as the member's default value may be outside the legal range for the enumeration, in which case an exception will occur during marshaling unless the member is explicitly set to a legal value.

To ensure that data members of primitive types are initialized to reasonable values, you can declare default values in your Slice definition. The default constructor initializes each of these data members to its declared value.

Exceptions also have a second constructor that has one parameter for each data member. This allows you to construct and initialize a class instance in a single statement (instead of first having to construct the instance and then assign to its members). For derived exceptions, this constructor has one parameter for each of the base class's data members, plus one parameter for each of the derived class's data members, in base-to-derived order.

C++ Mapping for Run-Time Exceptions

The Ice run time throws run-time exceptions for a number of pre-defined error conditions. All run-time exceptions directly or indirectly derive from `Ice::LocalException` (which, in turn, derives from `Ice::Exception`). `Ice::LocalException` has the usual member functions: `ice_name`, `ice_clone`, `ice_throw`, and (inherited from `Ice::Exception`), `ice_print`, `ice_file`, and `ice_line`.

Recall the inheritance diagram for user and run-time exceptions. By catching exceptions at the appropriate point in the hierarchy, you can handle exceptions according to the category of error they indicate:

- `Ice::Exception`
  This is the root of the complete inheritance tree. Catching `Ice::Exception` catches both user and run-time exceptions. As shown earlier, `Ice::Exception` is a typedef for `IceUtil::Exception`, and `IceUtil::Exception` inherits from `std::exception.`
- `Ice::UserException`
  This is the root exception for all user exceptions. Catching `Ice::UserException` catches all user exceptions (but not run-time exceptions).
- `Ice::LocalException`
  This is the root exception for all run-time exceptions. Catching `Ice::LocalException` catches all run-time exceptions (but not user exceptions).
- `Ice::TimeoutException`
  This is the base exception for both operation-invocation and connection-establishment timeouts.
- `Ice::ConnectTimeoutException`
  This exception is raised when the initial attempt to establish a connection to a server times out.

For example, a `ConnectTimeoutException` can be handled as `ConnectTimeoutException`, `TimeoutException`, `LocalExceptio n`, or `Exception`.

You will probably have little need to catch run-time exceptions as their most-derived type and instead catch them as `LocalException`; the fine-grained error handling offered by the remainder of the hierarchy is of interest mainly in the implementation of the Ice run time. Exceptions to this rule are the exceptions related to facet and object life cycles, which you may want to catch explicitly. These exceptions are `FacetNotExistException` and `ObjectNotExistException`, respectively.

See Also

- User Exceptions
- Run-Time Exceptions
- C++ Mapping for Identifiers
- C++ Mapping for Modules
- C++ Mapping for Built-In Types
- C++ Mapping for Enumerations
- C++ Mapping for Structures
- C++ Mapping for Sequences
- C++ Mapping for Dictionaries
- C++ Mapping for Constants
- C++ Mapping for Optional Values
- Versioning
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

C++ Mapping for Interfaces

The mapping of Slice interfaces revolves around the idea that, to invoke a remote operation, you call a member function on a local class instance that is a proxy for the remote object. This makes the mapping easy and intuitive to use because making a remote procedure call is no different from making a local procedure call (apart from error semantics).

On this page:

- Proxy Classes and Proxy Handles
  - Inheritance from Ice::Object
  - Proxy Handles
- Methods on Proxy Handles
  - Default constructor
  - Copy constructor
  - Assignment operator
  - Checked cast
  - Unchecked cast
  - Stream insertion and stringification
  - Static id
- Using Proxy Methods in C++
- Object Identity and Proxy Comparison in C++

Proxy Classes and Proxy Handles

On the client side, a Slice interface maps to a class with member functions that correspond to the operations on that interface. Consider the following simple interface:

Slice

```
module M {
    interface Simple {
        void op();
    };
};
```

The Slice compiler generates the following definitions for use by the client:

Copyright 2016, ZeroC, Inc.

C++

```
namespace IceProxy {
    namespace M {
        class Simple;
    }
}

namespace M {
    class Simple;
    typedef IceInternal::ProxyHandle< ::IceProxy::M::Simple> SimplePrx;
    typedef IceInternal::Handle< ::M::Simple> SimplePtr;
}

namespace IceProxy {
    namespace M {
        class Simple : public virtual IceProxy::Ice::Object {
        public:
            void op();
            void op(const Ice::Context&);
            // ...
        };
    };
}
```

As you can see, the compiler generates a *proxy class* `Simple` in the `IceProxy::M` namespace, as well as a *proxy* handle `M::SimplePrx`. In general, for a module `M`, the generated names are `::IceProxy::M::<`*`interface-name`*`>` and `::M::<`*`interface-name`*`>Prx`.

In the client's address space, an instance of `IceProxy::M::Simple` is the local ambassador for a remote instance of the `Simple` interface in a server and is known as a *proxy class instance*. All the details about the server-side object, such as its address, what protocol to use, and its object identity are encapsulated in that instance.

Inheritance from `Ice::Object`

`Simple` inherits from `IceProxy::Ice::Object`, reflecting the fact that all Ice interfaces implicitly inherit from `Ice::Object`. For each operation in the interface, the proxy class has two overloaded member functions of the same name. For the preceding example, we find that the operation `op` has been mapped to two member functions `op`.

One of the overloaded member functions has a trailing parameter of type `Ice::Context`. This parameter is for use by the Ice run time to store information about how to deliver a request; normally, you do not need to supply a value here and can pretend that the trailing parameter does not exist. (The parameter is also used by IceStorm.)

Proxy Handles

Client-side application code never manipulates proxy class instances directly. In fact, you are not allowed to instantiate a proxy class directly. The following code will not compile because the proxy class does not provide a public default constructor:

C++

```
IceProxy::M::Simple s;  // Compile-time error!
```

Proxy instances are always instantiated on behalf of the client by the Ice run time, so client code never has any need to instantiate a proxy directly. When the client receives a proxy from the run time, it is given a *proxy handle* to the proxy, of type `<`*`interface-name`*`>Prx` (`SimplePrx` for the preceding example). The client accesses the proxy via its proxy handle; the handle takes care of forwarding operation

Copyright 2016, ZeroC, Inc.

invocations to its underlying proxy, as well as reference-counting the proxy. This means that no memory-management issues can arise: deallocation of a proxy is automatic and happens once the last handle to the proxy disappears (goes out of scope).

Because the application code always uses proxy handles and never touches the proxy class directly, we usually use the term proxy to denote both proxy handle and proxy class. This reflects the fact that, in actual use, the proxy handle looks and feels like the underlying proxy class instance. If the distinction is important, we use the terms *proxy class*, *proxy class instance*, and *proxy handle*.

Methods on Proxy Handles

As we saw for the preceding example, the handle is actually a template of type `IceInternal::ProxyHandle` that takes the proxy class as the template parameter. This template has the usual default constructor, copy constructor, and assignment operator.

Default constructor

You can default-construct a proxy handle. The default constructor creates a proxy that points nowhere (that is, points at no object at all). If you invoke an operation on such a null proxy, you get an `IceUtil::NullHandleException`:

C++

```
try {
    SimplePrx s;        // Default-constructed proxy
    s->op();            // Call via nil proxy
    assert(0);          // Can't get here
} catch (const IceUtil::NullHandleException&) {
    cout << "As expected, got a NullHandleException" << endl;
}
```

Copy constructor

The copy constructor ensures that you can construct a proxy handle from another proxy handle. Internally, this increments a reference count on the proxy; the destructor decrements the reference count again and, once the count drops to zero, deallocates the underlying proxy class instance. That way, memory leaks are avoided:

C++

```
{                               // Enter new scope
    SimplePrx s1 = ...;         // Get a proxy from somewhere
    SimplePrx s2(s1);           // Copy-construct s2
    assert(s1 == s2);           // Assertion passes
}                               // Leave scope; s1, s2, and the
                                // underlying proxy instance
                                // are deallocated
```

Note the assertion in this example: proxy handles support comparison.

Assignment operator

You can freely assign proxy handles to each other. The handle implementation ensures that the appropriate memory-management activities take place. Self-assignment is safe and you do not have to guard against it:

Copyright 2016, ZeroC, Inc.

C++

```
SimplePrx s1 = ...;     // Get a proxy from somewhere
SimplePrx s2;           // s2 is nil
s2 = s1;                // both point at the same object
s1 = 0;                 // s1 is nil
s2 = 0;                 // s2 is nil
```

Widening assignments work implicitly. For example, if we have two interfaces, `Base` and `Derived`, we can widen a `DerivedPrx` to a `BasePrx` implicitly:

C++

```
BasePrx base;
DerivedPrx derived;
base = derived;         // Fine, no problem
derived = base;         // Compile-time error
```

Implicit narrowing conversions result in a compile error, so the usual C++ semantics are preserved: you can always assign a derived type to a base type, but not vice versa.

Checked cast

Proxy handles provide a `checkedCast` method:

C++

```
namespace IceInternal {
    template<typename T>
    class ProxyHandle : public IceUtil::HandleBase<T> {
    public:
        template<class Y>
        static ProxyHandle checkedCast(const ProxyHandle<Y>& r);

        template<class Y>

static ProxyHandle checkedCast(const ProxyHandle<Y>& r, const ::Ice::Context& c);

        // ...
    };
}
```

A checked cast has the same function for proxies as a C++ `dynamic_cast` has for pointers: it allows you to assign a base proxy to a derived proxy. If the type of the base proxy's target object is compatible with the derived proxy's static type, the assignment succeeds and, after the assignment, the derived proxy denotes the same object as the base proxy. Otherwise, if the type of the base proxy's target object is incompatible with the derived proxy's static type, the derived proxy is set to null. Here is an example to illustrate this:

Copyright 2016, ZeroC, Inc.

C++

```
BasePrx base = ...;     // Initialize base proxy
DerivedPrx derived = DerivedPrx::checkedCast(base);
if (derived) {
        // Base has run-time type Derived,
        // use derived...
} else {
        // Base has some other, unrelated type
}
```

The expression `DerivedPrx::checkedCast(base)` tests whether `base` points at an object of type `Derived` (or an object with a type that is derived from `Derived`). If so, the cast succeeds and `derived` is set to point at the same object as `base`. Otherwise, the cast fails and `derived` is set to the null proxy.

Note that `checkedCast` is a static member function so, to do a down-cast, you always use the syntax `<`*`interface-name`*`>Prx::checkedCast`.

Also note that you can use proxies in boolean contexts. For example, `if (proxy)` returns true if the proxy is not null.

A `checkedCast` typically results in a remote message to the server.The message effectively asks the server "is the object denoted by this reference of type Derived?"

Calling `checkedCast` on a proxy that is already of the desired proxy type returns immediately that proxy. Otherwise, `checkedCast` always calls `ice_isA` on the target object, and upon success, creates a new instance of the desired proxy class.

The reply from the server is communicated to the application code in form of a successful (non-null) or unsuccessful (null) result. Sending a remote message is necessary because, as a rule, there is no way for the client to find out what the actual run-time type of the target object is without confirmation from the server. (For example, the server may replace the implementation of the object for an existing proxy with a more derived one.) This means that you have to be prepared for a `checkedCast` to fail. For example, if the server is not running, you will receive a `ConnectFailedException;` if the server is running, but the object denoted by the proxy no longer exists, you will receive an `ObjectNotExistException`.

Unchecked cast

In some cases, it is known that an object supports a more derived interface than the static type of its proxy. For such cases, you can use an unchecked down-cast:

C++

```
namespace IceInternal {
    template<typename T>
    class ProxyHandle : public IceUtil::HandleBase<T> {
    public:
        template<class Y>
        static ProxyHandle uncheckedCast(const ProxyHandle<Y>& r);
        // ...
  };
}
```

An `uncheckedCast` provides a down-cast *without* consulting the server as to the actual run-time type of the object, for example:

Copyright 2016, ZeroC, Inc.

C++

```
BasePrx base = ...;     // Initialize to point at a Derived
DerivedPrx derived;
derived = DerivedPrx::uncheckedCast(base);
// Use derived...
```

You should use an `uncheckedCast` only if you are certain that the target object indeed supports the more derived type: an `uncheckedCast`, as the name implies, is not checked in any way; it does not contact the object in the server and, if it fails, it does not return null. If you use the proxy resulting from an incorrect `uncheckedCast` to invoke an operation, the behavior is undefined. Most likely, you will receive an `OperationNotExistException`, but, depending on the circumstances, the Ice run time may also report an exception indicating that unmarshaling has failed, or even silently return garbage results.

Calling `uncheckedCast` on a proxy that is already of the desired proxy type returns immediately that proxy. Otherwise, `uncheckedCast` creates a new instance of the desired proxy class.

Despite its dangers, `uncheckedCast` is still useful because it avoids the cost of sending a message to the server. And, particularly during initialization, it is common to receive a proxy of static type `Ice::Object`, but with a known run-time type. In such cases, an `uncheckedCast` saves the overhead of sending a remote message.

Stream insertion and stringification

For convenience, proxy handles also support insertion of a proxy into a stream, for example:

C++

```
Ice::ObjectPrx p = ...;
cout << p << endl;
```

This code is equivalent to writing:

C++

```
Ice::ObjectPrx p = ...;
cout << p->ice_toString() << endl;
```

Either code prints the stringified proxy. You could also achieve the same thing by writing:

C++

```
Ice::ObjectPrx p = ...;
cout << communicator->proxyToString(p) << endl;
```

The advantage of using the `ice_toString` member function instead of `proxyToString` is that you do not need to have the communicator available at the point of call.

Static id

The `ice_staticId` method returns the type ID of the proxy's Slice interface:

Copyright 2016, ZeroC, Inc.

C++

```
namespace IceInternal {
    template<typename T>
    class ProxyHandle : public IceUtil::HandleBase<T> {
    public:
        static const std::string& ice_staticId();
        // ...
  };
}
```

Here's an example that shows how to use it:

C++

```
BasePrx base = ...;
if(base->ice_id() == Derived::ice_staticId())
    // target object implements Derived
```

Applications normally use `checkedCast` instead of calling `ice_id` directly.

Using Proxy Methods in C++

The base proxy class `ObjectPrx` supports a variety of methods for customizing a proxy. Since proxies are immutable, each of these "factory methods" returns a copy of the original proxy that contains the desired modification. For example, you can obtain a proxy configured with a ten second timeout as shown below:

C++

```
Ice::ObjectPrx proxy = communicator->stringToProxy(...);
proxy = proxy->ice_timeout(10000);
```

A factory method returns a new proxy object if the requested modification differs from the current proxy, otherwise it returns the current proxy. With few exceptions, factory methods return a proxy of the same type as the current proxy, therefore it is generally not necessary to repeat a down-cast after using a factory method. The example below demonstrates these semantics:

C++

```
Ice::ObjectPrx base = communicator->stringToProxy(...);
HelloPrx hello = HelloPrx::checkedCast(base);
hello = hello->ice_timeout(10000); // Type is preserved
hello->sayHello();
```

The only exceptions are the factory methods `ice_facet` and `ice_identity`. Calls to either of these methods may produce a proxy for an object of an unrelated type, therefore they return a base proxy that you must subsequently down-cast to an appropriate type.

Object Identity and Proxy Comparison in C++

Copyright 2016, ZeroC, Inc.

Proxy handles support comparison using the following operators:

- `operator bool`
  Proxies have a conversion operator to `bool`. The operator returns true if a proxy is not null, and false otherwise. This allows you to write:

  C++

```
BasePrx base = ...;
if (base)
        // It's a non-nil proxy
if (!base)
        // It's a nil proxy
```

- `operator==`
  `operator!=`
  These operators permit you to compare proxies for equality and inequality. It's also legal to compare a proxy against the literal `0` to test whether a proxy is null, but we recommend using the `bool` operator instead.

- `operator<`
  `operator<=`
  `operator>`
  `operator>=`
  Proxies support comparison. This allows you to place proxies into STL containers such as maps or sorted lists.

Note that proxy comparison uses *all* of the information in a proxy for the comparison. This means that not only the object identity must match for a comparison to succeed, but other details inside the proxy, such as the protocol and endpoint information, must be the same. In other words, comparison with == and `!=` tests for *proxy* identity, *not* object identity. A common mistake is to write code along the following lines:

C++

```
Ice::ObjectPrx p1 = ...;        // Get a proxy...
Ice::ObjectPrx p2 = ...;        // Get another proxy...

if (p1 != p2) {
    // p1 and p2 denote different objects       // WRONG!
} else {
    // p1 and p2 denote the same object         // Correct
}
```

Even though `p1` and `p2` differ, they may denote the same Ice object. This can happen because, for example, both `p1` and `p2` embed the same object identity, but each use a different protocol to contact the target object. Similarly, the protocols may be the same, but denote different endpoints (because a single Ice object can be contacted via several different transport endpoints). In other words, if two proxies compare equal with ==, we know that the two proxies denote the same object (because they are identical in all respects); however, if two proxies compare unequal with ==, we know absolutely nothing: the proxies may or may not denote the same object.

To compare the object identities of two proxies, you can use helper functions in the `Ice` namespace:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {

    bool proxyIdentityLess(const ObjectPrx&, const ObjectPrx&);
    bool proxyIdentityEqual(const ObjectPrx&, const ObjectPrx&);
    bool proxyIdentityAndFacetLess(const ObjectPrx&, const ObjectPrx&);
    bool proxyIdentityAndFacetEqual(const ObjectPrx&, const ObjectPrx&);

}
```

The `proxyIdentityEqual` function returns true if the object identities embedded in two proxies are the same and ignores other information in the proxies, such as facet and transport information. To include the facet name in the comparison, use `proxyIdentityAndFacetEqual` instead.

The `proxyIdentityLess` function establishes a total ordering on proxies. It is provided mainly so you can use object identity comparison with STL sorted containers. (The function uses `name` as the major ordering criterion, and `category` as the minor ordering criterion.) The `proxyIdentityAndFacetLess` function behaves similarly to `proxyIdentityLess`, except that it also compares the facet names of the proxies when their identities are equal.

`proxyIdentityEqual` and `proxyIdentityAndFacetLess` allow you to correctly compare proxies for object identity. The example below demonstrates how to use `proxyIdentityEqual`:

C++

```
Ice::ObjectPrx p1 = ...;        // Get a proxy...
Ice::ObjectPrx p2 = ...;        // Get another proxy...

if (!Ice::proxyIdentityEqual(p1, p2) {
    // p1 and p2 denote different objects       // Correct
} else {
    // p1 and p2 denote the same object         // Correct
}
```

See Also

- Interfaces, Operations, and Exceptions
- Proxies for Ice Objects
- C++ Mapping for Operations
- Example of a File System Client in C++
- Versioning
- IceStorm

Copyright 2016, ZeroC, Inc.

C++ Mapping for Operations

On this page:

- Basic C++ Mapping for Operations
- Normal and idempotent Operations in C++
- Passing Parameters in C++
  - In-Parameters in C++
  - Out-Parameters in C++
  - Optional Parameters in C++
  - Chained Invocations in C++
- Exception Handling in C++
  - Exceptions and Out-Parameters in C++
  - Exceptions and Return Values in C++

Basic C++ Mapping for Operations

As we saw in the C++ mapping for interfaces, for each operation on an interface, the proxy class contains a corresponding member function with the same name. To invoke an operation, you call it via the proxy handle. For example, here is part of the definitions for our file system:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };
    // ...
};
```

The proxy class for the `Node` interface, tidied up to remove irrelevant detail, is as follows:

C++

```
namespace IceProxy {
    namespace Filesystem {
        class Node : virtual public IceProxy::Ice::Object {
        public:
            std::string name();
            // ...
        };
        typedef IceInternal::ProxyHandle<Node> NodePrx;
        // ...
    }
    // ...
}
```

The `name` operation returns a value of type `string`. Given a proxy to an object of type `Node`, the client can invoke the operation as follows:

Copyright 2016, ZeroC, Inc.

C++

```
NodePrx node = ...;             // Initialize proxy
string name = node->name();     // Get name via RPC
```

The proxy handle overloads `operator->` to forward method calls to the underlying proxy class instance which, in turn, sends the operation invocation to the server, waits until the operation is complete, and then unmarshals the return value and returns it to the caller.

Because the return value is of type `string`, it is safe to ignore the return value. For example, the following code contains no memory leak:

C++

```
NodePrx node = ...;             // Initialize proxy
node->name();                   // Useless, but no leak
```

This is true for all mapped Slice types: you can safely ignore the return value of an operation, no matter what its type — return values are always returned by value. If you ignore the return value, no memory leak occurs because the destructor of the returned value takes care of deallocating memory as needed.

Normal and `idempotent` Operations in C++

You can add an `idempotent` qualifier to a Slice operation. As far as the signature for the corresponding proxy methods is concerned, `idempotent` has no effect. For example, consider the following interface:

Slice

```
interface Example {
               string op1();
    idempotent string op2();
    idempotent void op3(string s);
};
```

The proxy class for this interface looks like this:

C++

```
namespace IceProxy {
    class Example : virtual public IceProxy::Ice::Object {
    public:
        std::string op1();
        std::string op2();              // idempotent
        void op3(const std::string&);   // idempotent
        // ...
    };
}
```


Copyright 2016, ZeroC, Inc.

Because `idempotent` affects an aspect of call dispatch, not interface, it makes sense for the mapping to be unaffected by the `idempotent` keyword.

Passing Parameters in C++

In-Parameters in C++

The parameter passing rules for the C++ mapping are very simple: parameters are passed either by value (for small values) or by `const` reference (for values that are larger than a machine word). Semantically, the two ways of passing parameters are identical: it is guaranteed that the value of a parameter will not be changed by the invocation.

Here is an interface with operations that pass parameters of various types from client to server:

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ClientToServer {
    void op1(int i, float f, bool b, string s);
    void op2(NumberAndString ns, StringSeq ss, StringTable st);
    void op3(ClientToServer* proxy);
};
```

The Slice compiler generates the following code for this definition:

Copyright 2016, ZeroC, Inc.

C++

```
struct NumberAndString {
    Ice::Int x;
    std::string str;
    // ...
};

typedef std::vector<std::string> StringSeq;

typedef std::map<Ice::Long, StringSeq> StringTable;

namespace IceProxy {
    class ClientToServer : virtual public IceProxy::Ice::Object {
    public:
        void op1(Ice::Int, Ice::Float, bool, const std::string&);
        void op2(const NumberAndString&, const StringSeq&, const StringTable&);
        void op3(const ClientToServerPrx&);
        // ...
    };
}
```

Given a proxy to a `ClientToServer` interface, the client code can pass parameters as in the following example:

C++

```
ClientToServerPrx p = ...;              // Get proxy...

p->op1(42, 3.14, true, "Hello world!"); // Pass simple literals

int i = 42;
float f = 3.14;
bool b = true;
string s = "Hello world!";
p->op1(i, f, b, s);                     // Pass simple variables

NumberAndString ns = { 42, "The Answer" };
StringSeq ss;
ss.push_back("Hello world!");
StringTable st;
st[0] = ss;
p->op2(ns, ss, st);                     // Pass complex variables

p->op3(p);                              // Pass proxy
```

You can pass either literals or variables to the various operations. Because everything is passed by value or `const` reference, there are no memory-management issues to consider.

Copyright 2016, ZeroC, Inc.

Out-Parameters in C++

The C++ mapping passes out-parameters by reference. Here is the Slice definition once more, modified to pass all parameters in the `out` direction:

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ServerToClient {
    void op1(out int i, out float f, out bool b, out string s);
    void op2(out NumberAndString ns, out StringSeq ss, out StringTable st);
    void op3(out ServerToClient* proxy);
};
```

The Slice compiler generates the following code for this definition:

C++

```
namespace IceProxy {
    class ServerToClient : virtual public IceProxy::Ice::Object {
    public:
       void op1(Ice::Int&, Ice::Float&, bool&, std::string&);
       void op2(NumberAndString&, StringSeq&, StringTable&);
       void op3(ServerToClientPrx&);
       // ...
    };
}
```

Given a proxy to a `ServerToClient` interface, the client code can pass parameters as in the following example:

Copyright 2016, ZeroC, Inc.

C++

```
ServerToClientPrx p = ...;      // Get proxy...

int i;
float f;
bool b;
string s;

p->op1(i, f, b, s);
// i, f, b, and s contain updated values now

NumberAndString ns;
StringSeq ss;
StringTable st;

p->op2(ns, ss, st);
// ns, ss, and st contain updated values now

p->op3(p);
// p has changed now!
```

Again, there are no surprises in this code: the caller simply passes variables to an operation; once the operation completes, the values of those variables will be set by the server.

It is worth having another look at the final call:

C++

```
p->op3(p);      // Weird, but well-defined
```

Here, p is the proxy that is used to dispatch the call. That same variable p is also passed as an out-parameter to the call, meaning that the server will set its value. In general, passing the same parameter as both an input and output parameter is safe: the Ice run time will correctly handle all locking and memory-management activities.

Optional Parameters in C++

The mapping for optional parameters is the same as for required parameters, except each optional parameter is encapsulated in an `IceUtil::Optional` value. Consider the following operation:

Slice

```
optional(1) int execute(optional(2) string params, out optional(3) float value);
```

The C++ mapping for this operation is shown below:

Copyright 2016, ZeroC, Inc.

C++

```
IceUtil::Optional<Ice::Int>
execute(const IceUtil::Optional<std::string>& params,
IceUtil::Optional<Ice::Float>& value, ...);
```

The constructors provided by the `IceUtil::Optional` template simplify the use of optional parameters:

C++

```
IceUtil::Optional<Ice::Int> i;
IceUtil::Optional<Ice::Float> v;

i = proxy->execute("--file log.txt", v); // string converted to Optional<string>
i = proxy->execute(IceUtil::None, v);    // params is unset

if(v)
    cout << "value = " << v.get() << endl;
```

For an optional output parameter, the Ice run time resets the client's `Optional` instance if the server does not supply a value for the parameter, therefore it is safe to pass an `Optional` instance that already has a value.

A well-behaved program must not assume that an optional parameter always has a value.

Chained Invocations in C++

Consider the following simple interface containing two operations, one to set a value and one to get it:

Slice

```
interface Name {
    string getName();
    void setName(string name);
};
```

Suppose we have two proxies to interfaces of type `Name`, `p1` and `p2`, and chain invocations as follows:

C++

```
p2->setName(p1->getName());
```

This works exactly as intended: the value returned by `p1` is transferred to `p2`. There are no memory-management or exception safety issues with this code.

Exception Handling in C++

Copyright 2016, ZeroC, Inc.

Any operation invocation may throw a run-time exception and, if the operation has an exception specification, may also throw user exceptions. Suppose we have the following simple interface:

Slice

```
exception Tantrum {
    string reason;
};

interface Child {
    void askToCleanUp() throws Tantrum;
};
```

Slice exceptions are thrown as C++ exceptions, so you can simply enclose one or more operation invocations in a `try-catch` block:

C++

```
ChildPrx child = ...;           // Get proxy...
try {
    child->askToCleanUp();      // Give it a try...
} catch (const Tantrum& t) {
    cout << "The child says: " << t.reason << endl;
}
```

Typically, you will catch only a few exceptions of specific interest around an operation invocation; other exceptions, such as unexpected run-time errors, will typically be dealt with by exception handlers higher in the hierarchy. For example:

Copyright 2016, ZeroC, Inc.

C++

```
void run()
{
    ChildPrx child = ...;       // Get proxy...
    try {
        child->askToCleanUp();  // Give it a try...
    } catch (const Tantrum& t) {
        cout << "The child says: " << t.reason << endl;

        child->scold();         // Recover from error...
    }
    child->praise();            // Give positive feedback...
}

int
main(int argc, char* argv[])
{
    int status = 1;
    try {
        // ...
        run();
        // ...
        status = 0;
    } catch (const Ice::Exception& e) {
        cerr << "Unexpected run?time error: " << e << endl;
    }
    // ...
    return status;
}
```

This code handles a specific exception of local interest at the point of call and deals with other exceptions generically. (This is also the strategy we used for our first simple application.)

For efficiency reasons, you should always catch exceptions by `const` reference. This permits the compiler to avoid calling the exception's copy constructor (and, of course, prevents the exception from being sliced to a base type).

Exceptions and Out-Parameters in C++

The Ice run time makes no guarantees about the state of out-parameters when an operation throws an exception: the parameter may have still have its original value or may have been changed by the operation's implementation in the target object. In other words, for out-parameters, Ice provides the weak exception guarantee [1] but does not provide the strong exception guarantee.

This is done for reasons of efficiency: providing the strong exception guarantee would require more overhead than can be justified.

Exceptions and Return Values in C++

For return values, C++ provides the guarantee that a variable receiving the return value of an operation will not be overwritten if an exception is thrown. (Of course, this guarantee holds only if you do not use the same variable as both an out-parameter and to receive the return value of an invocation).

See Also

- Operations
- Slice for a Simple File System

Copyright 2016, ZeroC, Inc.

- C++ Mapping for Interfaces

References

1. Sutter, H. 1999. Exceptional C++: 47 Engineering Puzzles, Programming Problems, and Solutions. Reading, MA: Addison-Wesley.

Copyright 2016, ZeroC, Inc.

C++ Mapping for Optional Values

On this page:

- The IceUtil::Optional Template
- The IceUtil::None Value

The `IceUtil::Optional` Template

The C++ mapping uses a template to hold the values of optional data members and parameters:

C++

```
namespace IceUtil {
template<typename T>
class Optional {
public:
    typedef T element_type;

    Optional();
    Optional(NoneType);
    Optional(const T&);

    template<typename Y>
    Optional(const Optional<Y>&);

    Optional(const Optional& r);

    Optional& operator=(NoneType);
    Optional& operator=(const T&);

    template<typename Y>
    Optional& operator=(const Optional<Y>&);
    Optional& operator=(const Optional&);

    const T& get() const;
    T& get();
    const T* operator->() const;
    T* operator->();
    const T& operator*() const;
    T& operator*();

    operator bool() const;
    bool operator!() const;

    void swap(Optional& other);

    ...
};
}
```

The `IceUtil::Optional` template provides constructors and assignment operators that allow you to initialize an instance using the element type or an existing optional value. The default constructor initializes an instance to an unset condition. The `get` method and dereference operators retrieve the current value held by the instance, or throw `IceUtil::OptionalNotSetException` if no value is

Copyright 2016, ZeroC, Inc.

currently set. Use the `bool` or `!` operators to test whether the instance has a value prior to dereferencing it. Finally, the `swap` method exchanges the state of two instances.

The `IceUtil::None` Value

The template includes a constructor and assignment operator that accept `NoneType`. Ice defines an instance of this type, `IceUtil::None`, that you can use to initialize (or reset) an `Optional` instance to an unset condition:

C++

```
IceUtil::Optional<int> i = 5;
i = IceUtil::None;
assert(!i); // true
```

You can pass `IceUtil::None` anywhere an `IceUtil::Optional` value is expected.

See Also

- Optional Data Members
- Operations
- C++ Mapping for Operations

Copyright 2016, ZeroC, Inc.

C++ Mapping for Classes

On this page:

- Basic C++ Mapping for Classes
- Inheritance from Ice::Object in C++
- Class Data Members in C++
- Class Constructors in C++
- Class Operations in C++
- Class Factories in C++

Basic C++ Mapping for Classes

A Slice class is mapped to a C++ class with the same name. The generated class contains a public data member for each Slice data member (just as for structures and exceptions), and a virtual member function for each operation. Consider the following class definition:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Slice compiler generates the following code for this definition:

Copyright 2016, ZeroC, Inc.

C++

```
class TimeOfDay;

typedef IceInternal::ProxyHandle<IceProxy::TimeOfDay> TimeOfDayPrx;
typedef IceInternal::Handle<TimeOfDay> TimeOfDayPtr;

class TimeOfDay : virtual public Ice::Object {
public:
    Ice::Short hour;
    Ice::Short minute;
    Ice::Short second;

    virtual std::string format() = 0;

    TimeOfDay() {};
    TimeOfDay(Ice::Short, Ice::Short, Ice::Short);

    virtual bool ice_isA(const std::string&);
    virtual const std::string& ice_id();
    static const std::string& ice_staticId();

    typedef TimeOfDayPrx ProxyType;
    typedef TimeOfDayPtr PointerType;

    // ...
};
```

The ProxyType and PointerType definitions are for template programming.

There are a number of things to note about the generated code:

1. The generated class `TimeOfDay` inherits from `Ice::Object`. This means that all classes implicitly inherit from `Ice::Object`, which is the ultimate ancestor of all classes. Note that `Ice::Object` is *not* the same as `IceProxy::Ice::Object`. In other words, you *cannot* pass a class where a proxy is expected and vice versa.
2. The generated class contains a public member for each Slice data member.
3. The generated class has a constructor that takes one argument for each data member, as well as a default constructor.
4. The generated class contains a pure virtual member function for each Slice operation.
5. The generated class contains additional member functions: `ice_isA`, `ice_id`, `ice_staticId`, and `ice_factory`.
6. The compiler generates a type definition `TimeOfDayPtr`. This type implements a smart pointer that wraps dynamically-allocated instances of the class. In general, the name of this type is <*class-name*>`Ptr`. Do not confuse this with <*class-name*>`Prx` — that type exists as well, but is the proxy handle for the class, not a smart pointer.

There is quite a bit to discuss here, so we will look at each item in turn.

Inheritance from `Ice::Object` in C++

Like interfaces, classes implicitly inherit from a common base class, `Ice::Object`. However, as shown in the figure below, classes inherited from `Ice::Object` instead of `Ice::ObjectPrx` (which is at the base of the inheritance hierarchy for proxies). As a result, you cannot pass a class where a proxy is expected (and vice versa) because the base types for classes and proxies are not compatible.

Copyright 2016, ZeroC, Inc.

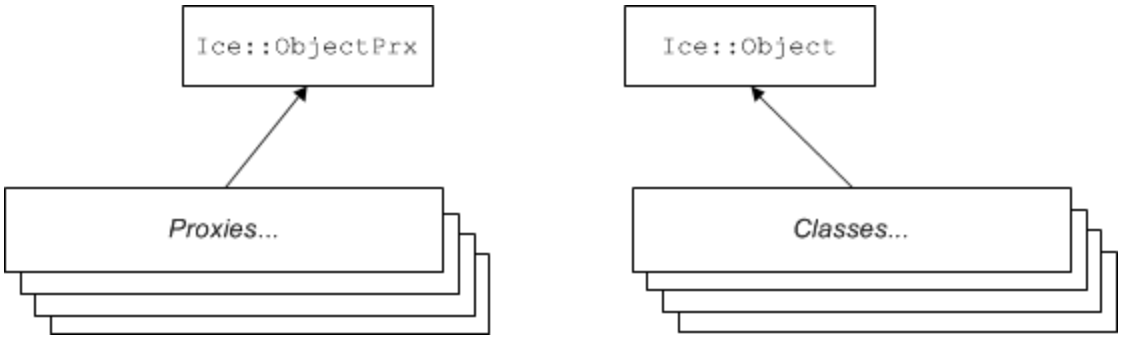

Inheritance from `Ice::ObjectPrx` and `Ice::Object`.

`Ice::Object` contains a number of member functions:

C++

```
namespace Ice {
    class Object : public virtual IceInternal::Shared {
    public:
        virtual bool ice_isA(const std::string&, const Current& = Current()) const;
        virtual void ice_ping(const Current&  = Current()) const;
        virtual std::vector<std::string> ice_ids(const Current& = Current()) const;
        virtual const std::string& ice_id(const Current& = Current()) const;
        static const std::string& ice_staticId();
        virtual ObjectPtr ice_clone() const;

        virtual void ice_preMarshal();
        virtual void ice_postUnmarshal();

        virtual void ice_collectable(bool);

        virtual DispatchStatus ice_dispatch(
            Ice::Request&,
            const DispatchInterceptorAsyncCallbackPtr& = 0);

        virtual bool operator==(const Object&) const;
        virtual bool operator!=(const Object&) const;
        virtual bool operator<(const Object&) const;
        virtual bool operator<=(const Object&) const;
        virtual bool operator>(const Object&) const;
        virtual bool operator>=(const Object&) const;
    };
}
```

The member functions of `Ice::Object` behave as follows:

- `ice_isA`
  This function returns `true` if the object supports the given type ID, and `false` otherwise.
- `ice_ping`
  As for interfaces, `ice_ping` provides a basic reachability test for the class. Note that `ice_ping` is normally only invoked on the proxy for a class that might be remote because a class instance that is local (in the caller's address space) can always be reached.
- `ice_ids`

Copyright 2016, ZeroC, Inc.

This function returns a string sequence representing all of the type IDs supported by this object, including `::Ice::Object`.

- `ice_id`
  This function returns the actual run-time type ID for a class. If you call `ice_id` through a smart pointer to a base instance, the returned type id is the actual (possibly more derived) type ID of the instance.
- `ice_staticId`
  This function returns the static type ID of a class.
- `ice_clone`
  This function makes a polymorphic shallow copy of a class.
- `ice_preMarshal`
  The Ice run time invokes this function prior to marshaling the object's state, providing the opportunity for a subclass to validate its declared data members.
- `ice_postUnmarshal`
  The Ice run time invokes this function after unmarshaling an object's state. A subclass typically overrides this function when it needs to perform additional initialization using the values of its declared data members.
- `ice_collectable`
  Determines whether this object, and by extension the graph of all objects reachable from this object, are eligible for garbage collection when all external references to the graph have been released.
- `ice_dispatch`
  This function dispatches an incoming request to a servant. It is used in the implementation of dispatch interceptors.
- `operator==`
  `operator!=`
  `operator<`
  `operator<=`
  `operator>`
  `operator>=`
  The comparison operators permit you to use classes as elements of STL sorted containers. Note that sort order, unlike for structures , is based on the memory address of the class, not on the contents of its data members of the class.

Class Data Members in C++

By default, data members of classes are mapped exactly as for structures and exceptions: for each data member in the Slice definition, the generated class contains a corresponding public data member. Optional data members are mapped to instances of the `IceUtil::Optional` template.

If you wish to restrict access to a data member, you can modify its visibility using the `protected` metadata directive. The presence of this directive causes the Slice compiler to generate the data member with protected visibility. As a result, the member can be accessed only by the class itself or by one of its subclasses. For example, the `TimeOfDay` class shown below has the `protected` metadata directive applied to each of its data members:

Slice

```
class TimeOfDay {
    ["protected"] short hour;   // 0 - 23
    ["protected"] short minute; // 0 - 59
    ["protected"] short second; // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Slice compiler produces the following generated code for this definition:

Copyright 2016, ZeroC, Inc.

C++

```
class TimeOfDay : virtual public Ice::Object {
public:

    virtual std::string format() = 0;

    // ...

protected:

    Ice::Short hour;
    Ice::Short minute;
    Ice::Short second;
};
```

For a class in which all of the data members are protected, the metadata directive can be applied to the class itself rather than to each member individually. For example, we can rewrite the `TimeOfDay` class as follows:

Slice

```
["protected"] class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

Class Constructors in C++

Classes have a default constructor that default-constructs each data member. Members having a complex type, such as strings, sequences, and dictionaries, are initialized by their own default constructor. However, the default constructor performs no initialization for members having one of the simple built-in types boolean, integer, floating point, or enumeration. For such a member, it is not safe to assume that the member has a reasonable default value. This is especially true for enumerated types as the member's default value may be outside the legal range for the enumeration, in which case an exception will occur during marshaling unless the member is explicitly set to a legal value.

To ensure that data members of primitive types are initialized to reasonable values, you can declare default values in your Slice definition. The default constructor initializes each of these data members to its declared value. Optional data members are unset unless they declare default values.

Classes also have a second constructor that has one parameter for each data member. This allows you to construct and initialize a class instance in a single statement. For each optional data member, its corresponding constructor parameter uses the same mapping as for operation parameters, allowing you to pass its initial value or `IceUtil::None` to indicate an unset value.

For derived classes, the constructor has one parameter for each of the base class's data members, plus one parameter for each of the derived class's data members, in base-to-derived order. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
class Base {
    int i;
};

class Derived extends Base {
    string s;
};
```

This generates:

C++

```
class Base : virtual public ::Ice::Object
{
public:
    ::Ice::Int i;

    Base() {};
    explicit Base(::Ice::Int);

    // ...
};

class Derived : public Base
{
public:
    ::std::string s;

    Derived() {};
    Derived(::Ice::Int, const ::std::string&);

    // ...
};
```

Note that single-parameter constructors are defined as `explicit`, to prevent implicit argument conversions.

By default, derived classes derive non-virtually from their base class. If you need virtual inheritance, you can enable it using the `["cpp:virtual"]` metadata directive.

Class Operations in C++

Operations of classes are mapped to pure virtual member functions in the generated class. This means that, if a class contains operations (such as the `format` operation of our `TimeOfDay` class), you must provide an implementation of the operation in a class that is derived from the generated class. For example:

Copyright 2016, ZeroC, Inc.

C++

```
class TimeOfDayI : virtual public TimeOfDay {
public:
    virtual std::string format() {
        std::ostringstream s;
        s << setw(2) << setfill('0') << hour << ":";
        s << setw(2) << setfill('0') << minute << ":";
        s << setw(2) << setfill('0') << second;
        return s.c_str();
    }

protected:
    virtual ~TimeOfDayI() {}  // Optional
};
```

We discuss the motivation for the protected destructor in Preventing Stack-Allocation of Class Instances.

Class Factories in C++

Having created a class such as `TimeOfDayI`, we have an implementation and we can instantiate the `TimeOfDayI` class, but we cannot receive it as the return value or as an out-parameter from an operation invocation. To see why, consider the following simple interface:

Slice

```
interface Time {
    TimeOfDay get();
};
```

When a client invokes the `get` operation, the Ice run time must instantiate and return an instance of the `TimeOfDay` class. However, `TimeOfDay` is an abstract class that cannot be instantiated. Unless we tell it, the Ice run time cannot magically know that we have created a `TimeOfDayI` class that implements the abstract `format` operation of the `TimeOfDay` abstract class. In other words, we must provide the Ice run time with a factory that knows that the `TimeOfDay` abstract class has a `TimeOfDayI` concrete implementation. The `Ice::Communicator` interface provides us with the necessary operations:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local interface ObjectFactory {
        Object create(string type);
        void destroy();
    };

    local interface Communicator {
        void addObjectFactory(ObjectFactory factory, string id);
        ObjectFactory findObjectFactory(string id);
        // ...
    };
};
```

To supply the Ice run time with a factory for our `TimeOfDayI` class, we must implement the `ObjectFactory` interface:

Slice

```
module Ice {
    local interface ObjectFactory {
        Object create(string type);
        void destroy();
    };
};
```

The object factory's `create` operation is called by the Ice run time when it needs to instantiate a `TimeOfDay` class. The factory's `destroy` operation is called by the Ice run time when its communicator is destroyed. A possible implementation of our object factory is:

C++

```
class ObjectFactory : public Ice::ObjectFactory {
public:
    virtual Ice::ObjectPtr create(const std::string& type) {
        assert(type == M::TimeOfDay::ice_staticId());
        return new TimeOfDayI;
    }
    virtual void destroy() {}
};
```

The `create` method is passed the type ID of the class to instantiate. For our `TimeOfDay` class, the type ID is `"::M::TimeOfDay"`. Our implementation of `create` checks the type ID: if it matches, the method instantiates and returns a `TimeOfDayI` object. For other type IDs, the method asserts because it does not know how to instantiate other types of objects.

Note that we used the `ice_staticId` method to obtain the type ID rather than embedding a literal string. Using a literal type ID string in your code is discouraged because it can lead to errors that are only detected at run time. For example, if a Slice class or one of its enclosing

Copyright 2016, ZeroC, Inc.

modules is renamed and the literal string is not changed accordingly, a receiver will fail to unmarshal the object and the Ice run time will raise `NoObjectFactoryException`. By using `ice_staticId` instead, we avoid any risk of a misspelled or obsolete type ID, and we can discover at compile time if a Slice class or module has been renamed.

Given a factory implementation, such as our `ObjectFactory`, we must inform the Ice run time of the existence of the factory:

C++

```
Ice::CommunicatorPtr ic = ...;
ic->addObjectFactory(new ObjectFactory, M::TimeOfDay::ice_staticId());
```

Now, whenever the Ice run time needs to instantiate a class with the type ID `"::M::TimeOfDay"`, it calls the `create` method of the registered `ObjectFactory` instance.

The `destroy` operation of the object factory is invoked by the Ice run time when the communicator is destroyed. This gives you a chance to clean up any resources that may be used by your factory. Do not call `destroy` on the factory while it is registered with the communicator — if you do, the Ice run time has no idea that this has happened and, depending on what your `destroy` implementation is doing, may cause undefined behavior when the Ice run time tries to next use the factory.

The run time guarantees that `destroy` will be the last call made on the factory, that is, `create` will not be called concurrently with `destroy`, and `create` will not be called once `destroy` has been called. However, calls to `create` can be made concurrently.

Note that you cannot register a factory for the same type ID twice: if you call `addObjectFactory` with a type ID for which a factory is registered, the Ice run time throws an `AlreadyRegisteredException`.

Finally, keep in mind that if a class has only data members, but no operations, you need not create and register an object factory to transmit instances of such a class. Only if a class has operations do you have to define and register an object factory.

See Also

- Classes
- Smart Pointers for Classes
- C++ Mapping for Operations
- C++ Mapping for Optional Values
- Asynchronous Method Invocation (AMI) in C++
- Dispatch Interceptors

Copyright 2016, ZeroC, Inc.

Smart Pointers for Classes

On this page:

- Automatic Memory Management with Smart Pointers
- Copying and Assignment of Classes
- Polymorphic Copying of Classes
- Null Smart Pointers
- Preventing Stack-Allocation of Class Instances
- Smart Pointers and Constructors
- Smart Pointers and Exception Safety
- Smart Pointers and Cycles
- Garbage Collection of Class Instances
- Smart Pointer Comparison

Automatic Memory Management with Smart Pointers

A recurring theme for C++ programmers is the need to deal with memory allocations and deallocations in their programs. The difficulty of doing so is well known: in the face of exceptions, multiple return paths from functions, and callee-allocated memory that must be deallocated by the caller, it can be extremely difficult to ensure that a program does not leak resources. This is particularly important in multi-threaded programs: if you do not rigorously track ownership of dynamic memory, a thread may delete memory that is still used by another thread, usually with disastrous consequences.

To alleviate this problem, Ice provides smart pointers for classes. These smart pointers use reference counting to keep track of each class instance and, when the last reference to a class instance disappears, automatically delete the instance.

Smart pointer classes are an example of the *RAII* (*Resource Acquisition Is Initialization*) idiom [1].

Smart pointers are generated by the Slice compiler for each class type. For a Slice class <*class-name*>, the compiler generates a C++ smart pointer called <*class-name*>`Ptr`. Rather than showing all the details of the generated class, here is the basic usage pattern: whenever you allocate a class instance on the heap, you simply assign the pointer returned from `new` to a smart pointer for the class. Thereafter, memory management is automatic and the class instance is deleted once the last smart pointer for it goes out of scope:

C++

```
{                                       // Open scope
    TimeOfDayPtr tod = new TimeOfDayI;  // Allocate instance
    // Initialize...
    tod->hour = 18;
    tod->minute = 11;
    tod->second = 15;
    // ...
}                                       // No memory leak here!
```

As you can see, you use `operator->` to access the members of the class via its smart pointer. When the `tod` smart pointer goes out of scope, its destructor runs and, in turn, the destructor takes care of calling `delete` on the underlying class instance, so no memory is leaked.

A smart pointer performs reference counting of its underlying class instance:

- The constructor of a class sets its reference count to zero.
- Initializing a smart pointer with a dynamically-allocated class instance causes the smart pointer to increment the reference count of the instance by one.
- Copy-constructing a smart pointer increments the reference count of the instance by one.
- Assigning one smart pointer to another increments the target's reference count and decrements the source's reference count. (Self-assignment is safe.)
- The destructor of a smart pointer decrements the reference count by one and calls `delete` on its class instance if the reference count drops to zero.

Suppose that we default-construct a smart pointer as follows:

Copyright 2016, ZeroC, Inc.

C++

```
TimeOfDayPtr tod;
```

This creates a smart pointer with an internal null pointer.

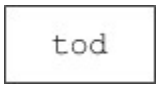

Newly initialized smart pointer.

Constructing a class instance creates that instance with a reference count of zero; the assignment to the smart pointer causes the smart pointer to increment the instance's reference count:

C++

```
tod = new TimeOfDayI;   // Refcount == 1
```

The resulting situation is shown below:

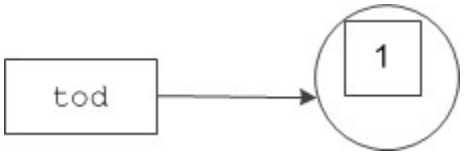

Initialized smart pointer.

Assigning or copy-constructing a smart pointer assigns and copy-constructs the smart pointer (not the underlying instance) and increments the reference count of the instance:

C++

```
TimeOfDayPtr tod2(tod); // Copy-construct tod2
TimeOfDayPtr tod3;
tod3 = tod;             // Assign to tod3
```

Here is the situation after executing these statements:

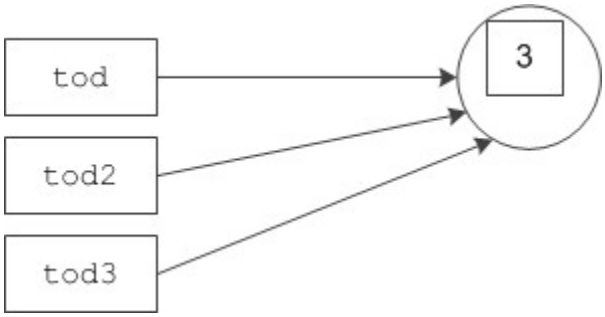

Three smart pointers pointing at the same class instance.

Copyright 2016, ZeroC, Inc.

Continuing the example, we can construct a second class instance and assign it to one of the original smart pointers, `tod2`:

C++

```
tod2 = new TimeOfDayI;
```

This decrements the reference count of the instance originally denoted by `tod2` and increments the reference count of the instance that is assigned to `tod2`. The resulting situation becomes the following:

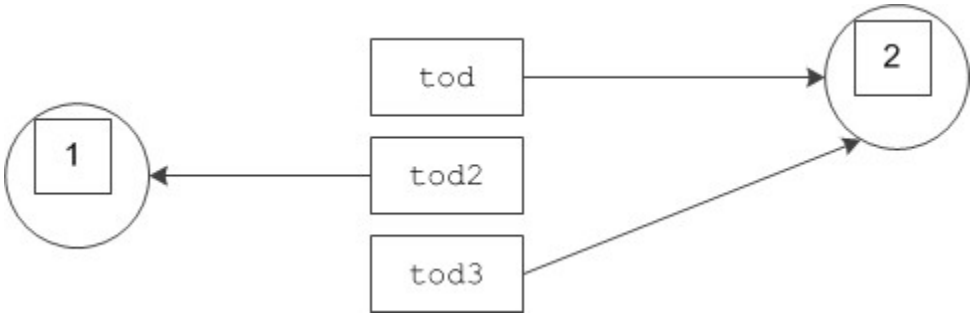

Three smart pointers and two instances.

You can clear a smart pointer by assigning zero to it:

C++

```
tod = 0;        // Clear handle
```

As you would expect, this decrements the reference count of the instance, as shown here:

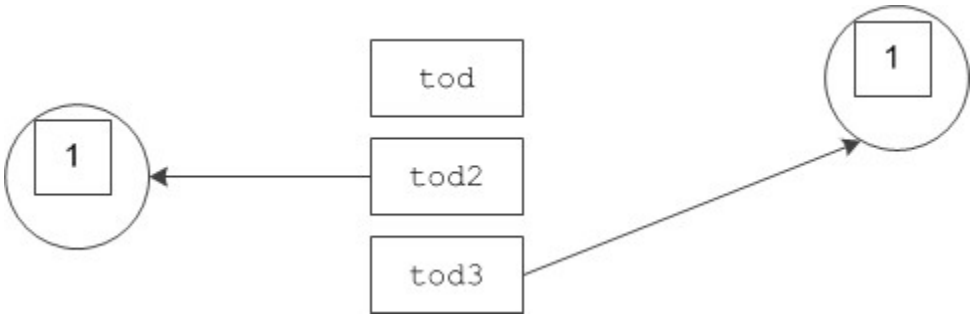

Decremented reference count after clearing a smart pointer.

If a smart pointer goes out of scope, is cleared, or has a new instance assigned to it, the smart pointer decrements the reference count of its instance; if the reference count drops to zero, the smart pointer calls `delete` to deallocate the instance. The following code snippet deallocates the instance on the right by assigning `tod2` to `tod3`:

C++

```
tod3 = tod2;
```

This results in the following situation:

Copyright 2016, ZeroC, Inc.

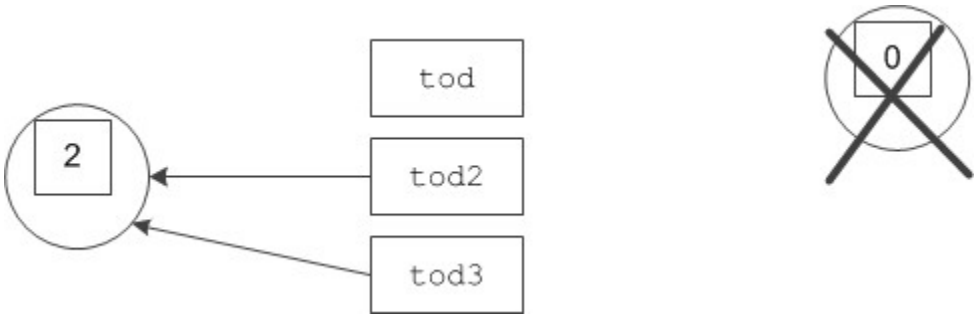

Deallocation of an instance with a reference count of zero.

Copying and Assignment of Classes

Classes have a default memberwise copy constructor and assignment operator, so you can copy and assign class instances:

C++

```
TimeOfDayPtr tod = new TimeOfDayI(2, 3, 4); // Create instance
TimeOfDayPtr tod2 = new TimeOfDayI(*tod);   // Copy instance

TimeOfDayPtr tod3 = new TimeOfDayI;
*tod3 = *tod;                               // Assign instance
```

Copying and assignment in this manner performs a memberwise shallow copy or assignment, that is, the source members are copied into the target members; if a class contains class members (which are mapped as smart pointers), what is copied or assigned is the smart pointer, not the target of the smart pointer.

To illustrate this, consider the following Slice definitions:

Slice

```
class Node {
    int i;
    Node next;
};
```

Assume that we initialize two instances of type `Node` as follows:

C++

```
NodePtr p1 = new Node(99, new Node(48, 0));
NodePtr p2 = new Node(23, 0);

// ...

*p2 = *p1; // Assignment
```


Copyright 2016, ZeroC, Inc.

After executing the first two statements, we have the situation shown below:

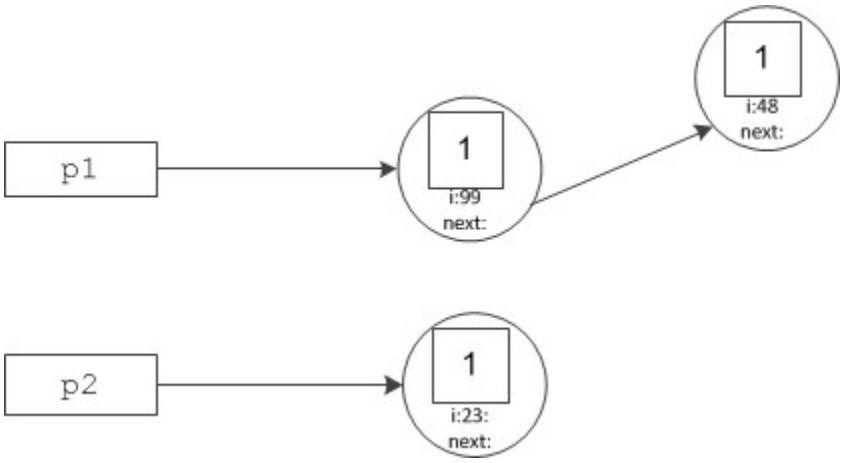

Class instances prior to assignment.

After executing the assignment statement, we end up with this result:

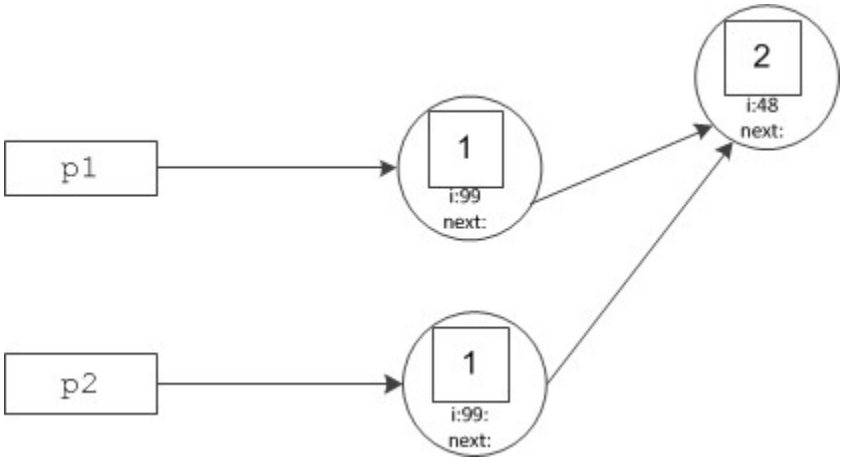

Class instances after assignment.

Note that copying and assignment also works for the implementation of abstract classes, such as our `TimeOfDayI` class, for example:

C++

```
class TimeOfDayI;

typedef IceUtil::Handle<TimeOfDayI> TimeOfDayIPtr;

class TimeOfDayI : virtual public TimeOfDay {
    // As before...
};
```

The default copy constructor and assignment operator will perform a memberwise copy or assignment of your implementation class:

Copyright 2016, ZeroC, Inc.

C++

```
TimeOfDayIPtr tod1 = new TimeOfDayI;
TimeOfDayIPtr tod2 = new TimeOfDayI(*tod1);     // Make copy
```

Of course, if your implementation class contains raw pointers (for which a memberwise copy would almost certainly be inappropriate), you must implement your own copy constructor and assignment operator that take the appropriate action (and probably call the base copy constructor or assignment operator to take care of the base part).

Note that the preceding code uses `TimeOfDayIPtr` as a typedef for `IceUtil::Handle<TimeOfDayI>`. This class is a template that contains the smart pointer implementation. If you want smart pointers for the implementation of an abstract class, you must define a smart pointer type as illustrated by this type definition.

Copying and assignment of classes also works correctly for derived classes: you can assign a derived class to a base class, but not vice-versa; during such an assignment, the derived part of the source class is sliced, as per the usual C++ assignment semantics.

Polymorphic Copying of Classes

As shown in Inheritance from `Ice::Object`, the C++ mapping generates an `ice_clone` member function for every class:

C++

```
class TimeOfDay : virtual public Ice::Object {
public:
    // ...

    virtual Ice::ObjectPtr ice_clone() const;
};
```

This member function makes a polymorphic shallow copy of a class: members that are not class members are deep copied; all members that are class members are shallow copied. To illustrate, consider the following class definition:

Slice

```
class Node {
    Node n1;
    Node n2;
};
```

Assume that we have an instance of this class, with the `n1` and `n2` members initialized to point at separate instances, as shown below:

Copyright 2016, ZeroC, Inc.

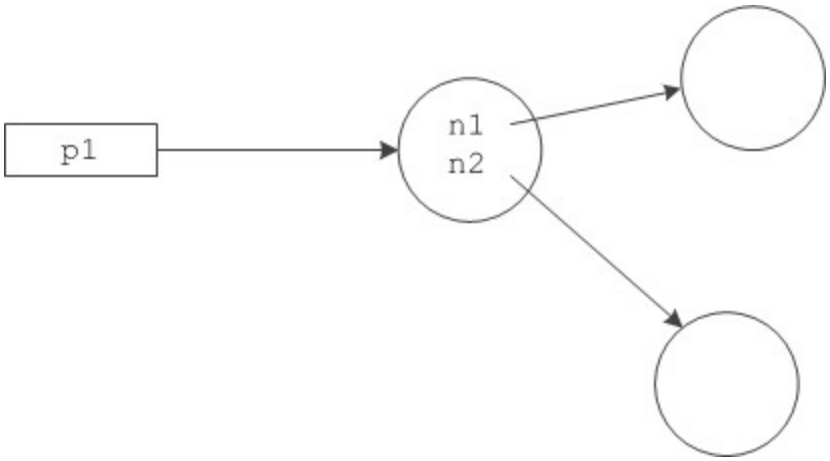

A class instance pointing at two other instances.

If we call `ice_clone` on the instance on the left, we end up with this situation:

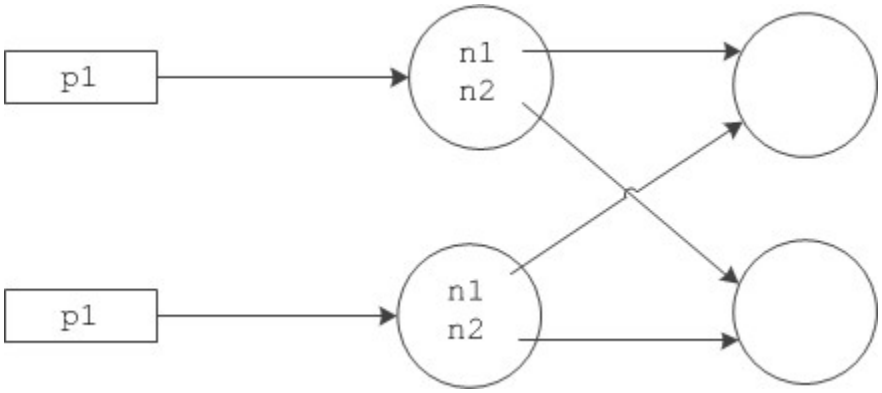

Resulting graph after calling `ice_clone` *on the left-most instance.*

As you can see, class members are shallow copied, that is, `ice_clone` makes a copy of the class instance on which it is invoked, but does not copy any class instances that are pointed at by the copied instance.

Note that `ice_clone` returns a value of type `Ice::ObjectPtr`, to avoid problems with compilers that do not support covariant return types. The generated `Ptr` classes contain a `dynamicCast` member that allows you to safely down-cast the return value of `ice_clone`. For example, the code to achieve the situation shown in the illustration above, looks as follows:

C++

```
NodePtr p1 = new Node(new Node, new Node);
NodePtr p2 = NodePtr::dynamicCast(p1->ice_clone());
```

`ice_clone` is generated by the Slice compiler for concrete classes (that is, classes that do not have operations). However, because classes with operations are abstract, the generated `ice_clone` for abstract classes cannot know how to instantiate an instance of the derived concrete class (because the name of the derived class is not known). This means that, for abstract classes, the generated `ice_clone` throws a `CloneNotImplementedException`.

If you want to clone the implementation of an abstract class, you must override the virtual `ice_clone` member in your concrete implementation class. For example:

Copyright 2016, ZeroC, Inc.

C++

```
class TimeOfDayI : public TimeOfDay {
public:
    virtual Ice::ObjectPtr ice_clone() const
    {
        return new TimeOfDayI(*this);
    }
};
```

Null Smart Pointers

A null smart pointer contains a null C++ pointer to its underlying instance. This means that if you attempt to dereference a null smart pointer, you get an `IceUtil::NullHandleException`:

C++

```
TimeOfDayPtr tod;               // Construct null handle

try {
    tod->minute = 0;            // Dereference null handle
    assert(false);              // Cannot get here
} catch (const IceUtil::NullHandleException&) {
    ; // OK, expected
} catch (...) {
    assert(false);              // Must get NullHandleException
}
```

Preventing Stack-Allocation of Class Instances

The Ice C++ mapping expects class instances to be allocated on the heap. Allocating class instances on the stack or in static variables is pragmatically useless because all the Ice APIs expect parameters that are smart pointers, not class instances. This means that, to do anything with a stack-allocated class instance, you must initialize a smart pointer for the instance. However, doing so does not work because it inevitably leads to a crash:

C++

```
{                               // Enter scope
    TimeOfDayI t;               // Stack-allocated class instance
    TimeOfDayPtr todp;          // Handle for a TimeOfDay instance

    todp = &t;                  // Legal, but dangerous
    // ...
}                               // Leave scope, looming crash!
```

This goes badly wrong because, when `todp` goes out of scope, it decrements the reference count of the class to zero, which then calls `dele`

Copyright 2016, ZeroC, Inc.

te on itself. However, the instance is stack-allocated and cannot be deleted, and we end up with undefined behavior (typically, a core dump).

The following attempt to fix this is also doomed to failure:

C++

```
{                               // Enter scope
    TimeOfDayI t;               // Stack-allocated class instance
    TimeOfDayPtr todp;          // Handle for a TimeOfDay instance

    todp = &t;                  // Legal, but dangerous
    // ...
    todp = 0;                   // Crash imminent!
}
```

This code attempts to circumvent the problem by clearing the smart pointer explicitly. However, doing so also causes the smart pointer to drop the reference count on the class to zero, so this code ends up with the same call to `delete` on the stack-allocated instance as the previous example.

The upshot of all this is: **never allocate a class instance on the stack or in a static variable**. The C++ mapping assumes that all class instances are allocated on the heap and no amount of coding trickery will change this.

You could abuse the `__setNoDelete` member to disable deallocation, but we strongly discourage you from doing this.

You can prevent allocation of class instances on the stack or in static variables by adding a protected destructor to your implementation of the class: if a class has a protected destructor, allocation must be made with `new`, and static or stack allocation causes a compile-time error. In addition, explicit calls to `delete` on a heap-allocated instance also are flagged at compile time.

Tip
You may want to habitually add a protected destructor to your implementation of abstract Slice classes to protect yourself from accidental heap allocation, as shown in Class Operations. (For Slice classes that do not have operations, the Slice compiler automatically adds a protected destructor.)

Smart Pointers and Constructors

Slice classes inherit their reference-counting behavior from the `IceUtil::Shared` class, which ensures that reference counts are managed in a thread-safe manner. When a stack-allocated smart pointer goes out of scope, the smart pointer invokes the `__decRef` function on the reference-counted object. Ignoring thread-safety issues, `__decRef` is implemented like this:

C++

```
void
IceUtil::Shared::__decRef()
{
    if (--_ref == 0 && !_noDelete)
        delete this;
}
```

In other words, when the smart pointer calls `__decRef` on the pointed-at instance and the reference count reaches zero (which happens when the last smart pointer for a class instance goes out of scope), the instance self-destructs by calling `delete this`.

However, as you can see, the instance self-destructs only if `_noDelete` is false (which it is by default, because the constructor initializes it

Copyright 2016, ZeroC, Inc.

to false). You can call `__setNoDelete(true)` to prevent this self-destruction and, later, call `__setNoDelete(false)` to enable it again. This is necessary if, for example, a class in its constructor needs to pass `this` to another function:

C++

```
void someFunction(const TimeOfDayPtr& t)
{
    // ...
}

TimeOfDayI::TimeOfDayI()
{
    someFunction(this); // Trouble looming here!
}
```

At first glance, this code looks innocuous enough. While `TimeOfDayI` is being constructed, it passes `this` to `someFunction`, which expects a smart pointer. The compiler constructs a temporary smart pointer at the point of call (because the smart pointer template has a single-argument constructor that accepts a pointer to a heap-allocated instance, so the constructor acts as a conversion function). However, this code fails badly. The `TimeOfDayI` instance is constructed with a statement such as:

C++

```
TimeOfDayPtr tod = new TimeOfDayI;
```

The constructor of `TimeOfDayI` is called by `operator new` and, when the constructor starts executing, the reference count of the instance is zero (because that is what the reference count is initialized to by the `Shared` base class of `TimeOfDayI`). When the constructor calls `someFunction`, the compiler creates a temporary smart pointer, which increments the reference count of the instance and, once `someFunction` completes, the compiler dutifully destroys that temporary smart pointer again. But, of course, that drops the reference count back to zero and causes the `TimeOfDayI` instance to self-destruct by calling `delete this`. The net effect is that the call to `new TimeOfDayI` returns a pointer to an already deleted object, which is likely to cause the program to crash.

To get around the problem, you can call `__setNoDelete`:

C++

```
TimeOfDayI::TimeOfDayI()
{
    __setNoDelete(true);
    someFunction(this);
    __setNoDelete(false);
}
```

The code disables self-destruction while `someFunction` uses its temporary smart pointer by calling `__setNoDelete(true)`. By doing this, the reference count of the instance is incremented before `someFunction` is called and decremented back to zero when `someFunction` completes without causing the object to self-destruct. The constructor then re-enables self-destruction by calling `__setNoDelete(false)` before returning, so the statement

Copyright 2016, ZeroC, Inc.

C++

```
TimeOfDayPtr tod = new TimeOfDayI;
```

does the usual thing, namely to increment the reference count of the object to 1, despite the fact that a temporary smart pointer existed while the constructor ran.

In general, whenever a class constructor passes `this` to a function or another class that accepts a smart pointer, you must temporarily disable self-destruction.

Smart Pointers and Exception Safety

Smart pointers are exception safe: if an exception causes the thread of execution to leave a scope containing a stack-allocated smart pointer, the C++ run time ensures that the smart pointer's destructor is called, so no resource leaks can occur:

C++

```
{ // Enter scope...

    TimeOfDayPtr tod = new TimeOfDayI; // Allocate instance

    someFuncThatMightThrow();          // Might throw...

    // ...

} // No leak here, even if an exception is thrown
```

If an exception is thrown, the destructor of `tod` runs and ensures that it deallocates the underlying class instance.

There is one potential pitfall you must be aware of though: if a constructor of a base class throws an exception, and another class instance holds a smart pointer to the instance being constructed, you can end up with a double deallocation. You can use the `__setNoDelete` mech anism to temporarily disable self-destruction in this case, as described above.

Smart Pointers and Cycles

One thing you need to be aware of is the inability of reference counting to deal with cyclic dependencies. For example, consider the following simple self-referential class:

Slice

```
class Node {
    int val;
    Node next;
};
```

Intuitively, this class implements a linked list of nodes. As long as there are no cycles in the list of nodes, everything is fine, and our smart pointers will correctly deallocate the class instances. However, if we introduce a cycle, we have a problem:

Copyright 2016, ZeroC, Inc.

C++

```
{                          // Open scope...

    NodePtr n1 = new Node; // N1 refcount == 1
    NodePtr n2 = new Node; // N2 refcount == 1
    n1->next = n2;         // N2 refcount == 2
    n2->next = n1;         // N1 refcount == 2

} // Destructors run:      // N2 refcount == 1,
                           // N1 refcount == 1, memory leak!
```

The nodes pointed to by `n1` and `n2` do not have names but, for the sake of illustration, let us assume that `n1`'s node is called N1, and `n2`'s node is called N2. When we allocate the N1 instance and assign it to `n1`, the smart pointer `n1` increments N1's reference count to 1. Similarly, N2's reference count is 1 after allocating the node and assigning it to `n2`. The next two statements set up a cyclic dependency between `n1` and `n2` by making their `next` pointers point at each other. This sets the reference count of both N1 and N2 to 2. When the enclosing scope closes, the destructor of `n2` is called first and decrements N2's reference count to 1, followed by the destructor of `n1`, which decrements N1's reference count to 1. The net effect is that neither reference count ever drops to zero, so both N1 and N2 are leaked.

Garbage Collection of Class Instances

The previous example illustrates a problem that is generic to using reference counts for deallocation: if a cyclic dependency exists anywhere in a graph (possibly via many intermediate nodes), all nodes in the cycle are leaked.

To avoid memory leaks due to such cycles, Ice for C++ includes a garbage collection facility. The facility identifies class instances that are part of a cycle but are no longer reachable from the program and deletes such instances. Applications must assist the Ice run time in this process by indicating when a graph is eligible for collection. For eligible graphs, Ice makes a sweep of the graph each time a reference count to one of the graph's objects is decremented.

Two components of the garbage collection facility influence its behavior:

- The `Ice.CollectObjects` property determines whether Ice assumes all object graphs are eligible for collection by default.
- The `Object::ice_collectable(bool)` method allows an application to indicate whether an object (and by extension, the graph of objects reachable via this object) is eligible for collection.

The correct operation of the garbage collection facility relies on the assumption that all eligible object graphs are immutable. If an application needs to make changes that could affect the structure of the graph, it must disable collection for that graph by calling `ice_collectable(false)` on any root object in the graph. Once the changes are complete, call `ice_collectable(true)` on any root object in the graph to make it eligible again. Modifying the structure of an eligible graph has undefined behavior.

In general, there are two strategies you can use for garbage collection depending on your application's requirements:

1. Set `Ice.CollectObjects=1` so that Ice assumes all object graphs are eligible for collection by default. This is recommended for applications that receive object graphs but rarely modify them. For those situations where an application needs to modify a graph, surround the modification with calls to `ice_collectable` as shown below:

Copyright 2016, ZeroC, Inc.

C++

```
NodePtr graph = proxy->getGraph(); // Eligible by default
graph->ice_collectable(false);
// modify graph...
graph->ice_collectable(true);
graph = 0; // Starts a sweep
```

2. Set `Ice.CollectObjects=0` (the default setting) so that Ice does not attempt to collect object graphs except for those explicitly marked by the application. Use this setting for applications that typically modify the structure of the graphs they receive. Call `ice_collectable(true)` to mark a graph as eligible:

C++

```
NodePtr graph = proxy->getGraph(); // Ineligible by default
// modify graph...
graph->ice_collectable(true);
graph = 0; // Starts a sweep
```

As mentioned earlier, an application does not take any action to force a collection; rather, the collection occurs automatically when a reference count is decremented.

To minimize overhead, GC-related behavior is only enabled for those Slice classes whose data members can refer to Ice objects. Furthermore, graph traversal only occurs for those objects that are part of a cycle and marked as collectable.

Smart Pointer Comparison

As for proxy handles, class handles support the comparison operators ==, `!=`, and <. This allows you to use class handles in STL sorted containers. Be aware that, for smart pointers, object identity is not used for the comparison, because class instances do not have identity. Instead, these operators simply compare the memory address of the classes they point to. This means that `operator==` returns true only if two smart pointers point at the same physical class instance:

Copyright 2016, ZeroC, Inc.

C++

```
// Create a class instance and initialize
//
TimeOfDayIPtr p1 = new TimeOfDayI;
p1->hour = 23;
p1->minute = 10;
p1->second = 18;

// Create another class instance with
// the same member values
//
TimeOfDayIPtr p2 = new TimeOfDayI;
p2->hour = 23;
p2->minute = 10;
p2->second = 18;

assert(p1 != p2);       // The two do not compare equal

TimeOfDayIPtr p3 = p1;  // Point at first class again

assert(p1 == p3);       // Now they compare equal
```

See Also

- Classes
- C++ Mapping for Classes
- Asynchronous Method Invocation (AMI) in C++
- The Server-Side main Function in C++
- Properties and Configuration
- The C++ Shared and SimpleShared Classes

References

1. Stroustrup, B. 1997. *The C++ Programming Language*. Reading, MA: Addison-Wesley.

Copyright 2016, ZeroC, Inc.

Asynchronous Method Invocation (AMI) in C++

Asynchronous Method Invocation (AMI) is the term used to describe the client-side support for the asynchronous programming model. AMI supports both oneway and twoway requests, but unlike their synchronous counterparts, AMI requests never block the calling thread. When a client issues an AMI request, the Ice run time hands the message off to the local transport buffer or, if the buffer is currently full, queues the request for later delivery. The application can then continue its activities and poll or wait for completion of the invocation, or receive a callback when the invocation completes.

AMI is transparent to the server: there is no way for the server to tell whether a client sent a request synchronously or asynchronously.

On this page:

- Basic Asynchronous API in C++
  - Asynchronous Proxy Methods in C++
  - Asynchronous Exception Semantics in C++
- AsyncResult Class in C++
- Polling for Completion in C++
- Generic Completion Callbacks in C++
  - Using Cookies for Generic Completion Callbacks in C++
- Type-Safe Completion Callbacks in C++
  - Using Cookies for Type-Safe Completion Callbacks in C++
- Lambda Completion Callbacks in C++
- Asynchronous Oneway Invocations in C++
- Flow Control in C++
- Asynchronous Batch Requests in C++
- Concurrency Semantics for AMI in C++

Basic Asynchronous API in C++

Consider the following simple Slice definition:

Slice

```
module Demo {
    interface Employees {
        string getName(int number);
    };
};
```

Asynchronous Proxy Methods in C++

Besides the synchronous proxy methods, `slice2cpp` generates the following asynchronous proxy methods:

C++

```
Ice::AsyncResultPtr begin_getName(Ice::Int number);
Ice::AsyncResultPtr begin_getName(Ice::Int number, const Ice::Context& __ctx)

std::string end_getName(const Ice::AsyncResultPtr&);
```

Four additional overloads of `begin_getName` are generated for use with generic callbacks and type-safe callbacks.

As you can see, the single `getName` operation results in `begin_getName` and `end_getName` methods. (The `begin_` method is overloaded so you can pass a per-invocation context.)

- The `begin_getName` method sends (or queues) an invocation of `getName`. This method does not block the calling thread.
- The `end_getName` method collects the result of the asynchronous invocation. If, at the time the calling thread calls `end_getName`,

Copyright 2016, ZeroC, Inc.

the result is not yet available, the calling thread blocks until the invocation completes. Otherwise, if the invocation completed some time before the call to `end_getName`, the method returns immediately with the result.

A client could call these methods as follows:

C++

```
EmployeesPrx e = ...;
Ice::AsyncResultPtr r = e->begin_getName(99);

// Continue to do other things here...

string name = e->end_getName(r);
```

Because `begin_getName` does not block, the calling thread can do other things while the operation is in progress.

Note that `begin_getName` returns a value of type `AsyncResultPtr`. The `AsyncResult` associated with this smart pointer contains the state that the Ice run time requires to keep track of the asynchronous invocation. You must pass the `AsyncResultPtr` that is returned by the `begin_` method to the corresponding `end_` method.

The `begin_` method has one parameter for each in-parameter of the corresponding Slice operation. Similarly, the `end_` method has one out-parameter for each out-parameter of the corresponding Slice operation (plus the `AsyncResultPtr` parameter). For example, consider the following operation:

Slice

```
double op(int inp1, string inp2, out bool outp1, out long outp2);
```

The `begin_op` and `end_op` methods have the following signature:

C++

```
Ice::AsyncResultPtr begin_op(Ice::Int inp1, const ::std::string& inp2)

Ice::Double end_op(bool& outp1, Ice::Long& outp2, const Ice::AsyncResultPtr&);
```

Asynchronous Exception Semantics in C++

If an invocation raises an exception, the exception is thrown by the `end_` method, even if the actual error condition for the exception was encountered during the `begin_` method ("on the way out"). The advantage of this behavior is that all exception handling is located with the code that calls the `end_` method (instead of being present twice, once where the `begin_` method is called, and again where the `end_` method is called).

There is one exception to the above rule: if you destroy the communicator and then make an asynchronous invocation, the `begin_` method throws `CommunicatorDestroyedException`. This is necessary because, once the run time is finalized, it can no longer throw an exception from the `end_` method.

The only other exception that is thrown by the `begin_` and `end_` methods is `IceUtil::IllegalArgumentException`. This exception indicates that you have used the API incorrectly. For example, the `begin_` method throws this exception if you call an operation that has a return value or out-parameters on a oneway proxy. Similarly, the `end_` method throws this exception if you use a different proxy to call the `end_` method than the proxy you used to call the `begin_` method, or if the `AsyncResult` you pass to the `end_` method was obtained by calling the `begin_` method for a different operation.

`AsyncResult` Class in C++

Copyright 2016, ZeroC, Inc.

The `AsyncResult` that is returned by the `begin_` method encapsulates the state of the asynchronous invocation:

C++

```
class AsyncResult : public Ice::LocalObject, private IceUtil::noncopyable {
public:
    bool operator==(const AsyncResult&) const;
    bool operator<(const AsyncResult&) const;

    Int getHash() const;

    void cancel();

    CommunicatorPtr getCommunicator() const;
    ConnectionPtr getConnection() const;
    ObjectPrx getProxy() const;
    const string& getOperation() const;
    LocalObjectPtr getCookie() const;

    bool isCompleted() const;
    void waitForCompleted();

    bool isSent() const;
    void waitForSent();

    void throwLocalException() const;

    bool sentSynchronously() const;
};
```

The methods have the following semantics:

- `bool operator==(const AsyncResult&) const`
  `bool operator<(const AsyncResult&) const`
  `Int getHash() const`
  These methods allow you to create ordered or hashed collections of pending asynchronous invocations. This is useful, for example, if you can have a number of outstanding requests, and need to pass state between the `begin_` and the `end_` methods. In this case, you can use the returned `AsyncResult` objects as keys into a map that stores the state for each call.
- `void cancel()`
  This method prevents a queued invocation from being sent or, if the invocation has already been sent, ignores a reply if the server sends one. `cancel` is a local operation and has no effect on the server. A canceled invocation is considered to be completed, meaning `isCompleted` returns true, and the result of the invocation is an `Ice::InvocationCanceledException`.
- `CommunicatorPtr getCommunicator() const`
  This method returns the communicator that sent the invocation.
- `virtual ConnectionPtr getConnection() const`
  This method returns the connection that was used for the invocation. Note that, for typical asynchronous proxy invocations, this method returns a nil value because the possibility of automatic retries means the connection that is currently in use could change unexpectedly. The `getConnection` method only returns a non-nil value when the `AsyncResult` object is obtained by calling `beg in_flushBatchRequests` on a `Connection` object.
- `virtual ObjectPrx getProxy() const`
  This method returns the proxy that was used to call the `begin_` method, or nil if the `AsyncResult` object was not obtained via an asynchronous proxy invocation.
- `const string& getOperation() const`
  This method returns the name of the operation.

Copyright 2016, ZeroC, Inc.

- `LocalObjectPtr getCookie() const`
  This method returns the cookie that was passed to the `begin_` method. If you did not pass a cookie to the `begin_` method, the return value is null.
- `bool isCompleted() const`
  This method returns true if, at the time it is called, the result of an invocation is available, indicating that a call to the `end_` method will not block the caller. Otherwise, if the result is not yet available, the method returns false.
- `void waitForCompleted()`
  This method blocks the caller until the result of an invocation becomes available.
- `bool isSent() const`
  When you call the `begin_` method, the Ice run time attempts to write the corresponding request to the client-side transport. If the transport cannot accept the request, the Ice run time queues the request for later transmission. `isSent` returns true if, at the time it is called, the request has been written to the local transport (whether it was initially queued or not). Otherwise, if the request is still queued or an exception occurred before the request could be sent, `isSent` returns false.
- `void waitForSent()`
  This method blocks the calling thread until a request has been written to the client-side transport, or an exception occurs. After `waitForSent` returns, `isSent` returns true if the request was successfully written to the client-side transport, or false if an exception occurred. In the case of a failure, you can call the corresponding `end_` method or `throwLocalException` to obtain the exception.
- `void throwLocalException() const`
  This method throws the local exception that caused the invocation to fail. If no exception has occurred yet, `throwLocalException` does nothing.
- `bool sentSynchronously() const`
  This method returns true if a request was written to the client-side transport without first being queued. If the request was initially queued, `sentSynchronously` returns false (independent of whether the request is still in the queue or has since been written to the client-side transport).

Polling for Completion in C++

The `AsyncResult` methods allow you to poll for call completion. Polling is useful in a variety of cases. As an example, consider the following simple interface to transfer files from client to server:

Slice

```
interface FileTransfer
{
    void send(int offset, ByteSeq bytes);
};
```

The client repeatedly calls `send` to send a chunk of the file, indicating at which offset in the file the chunk belongs. A naïve way to transmit a file would be along the following lines:

Copyright 2016, ZeroC, Inc.

C++

```
FileHandle file = open(...);
FileTransferPrx ft = ...;
const int chunkSize = ...;

Ice::Int offset = 0;
while (!file.eof()) {
    ByteSeq bs;
    bs = file.read(chunkSize); // Read a chunk
    ft->send(offset, bs);      // Send the chunk
    offset += bs.size();
}
```

This works, but not very well: because the client makes synchronous calls, it writes each chunk on the wire and then waits for the server to receive the data, process it, and return a reply before writing the next chunk. This means that both client and server spend much of their time doing nothing — the client does nothing while the server processes the data, and the server does nothing while it waits for the client to send the next chunk.

Using asynchronous calls, we can improve on this considerably:

Copyright 2016, ZeroC, Inc.

C++

```
FileHandle file = open(...);
FileTransferPrx ft = ...;
const int chunkSize = ...;
Ice::Int offset = 0;

list<Ice::AsyncResultPtr> results;
const int numRequests = 5;

while (!file.eof()) {
    ByteSeq bs;
    bs = file.read(chunkSize);

    // Send up to numRequests + 1 chunks asynchronously.
    Ice::AsyncResultPtr r = ft->begin_send(offset, bs);
    offset += bs.size();

    // Wait until this request has been passed to the transport.
    r->waitForSent();
    results.push_back(r);

    // Once there are more than numRequests, wait for the least
    // recent one to complete.
    while (results.size() > numRequests) {
        Ice::AsyncResultPtr r = results.front();
        results.pop_front();
        r->waitForCompleted();
    }
}

// Wait for any remaining requests to complete.
while (!results.empty()) {
    Ice::AsyncResultPtr r = results.front();
    results.pop_front();
    r->waitForCompleted();
}
```

With this code, the client sends up to `numRequests + 1` chunks before it waits for the least recent one of these requests to complete. In other words, the client sends the next request without waiting for the preceding request to complete, up to the limit set by `numRequests`. In effect, this allows the client to "keep the pipe to the server full of data": the client keeps sending data, so both client and server continuously do work.

Obviously, the correct chunk size and value of `numRequests` depend on the bandwidth of the network as well as the amount of time taken by the server to process each request. However, with a little testing, you can quickly zoom in on the point where making the requests larger or queuing more requests no longer improves performance. With this technique, you can realize the full bandwidth of the link to within a percent or two of the theoretical bandwidth limit of a native socket connection.

Generic Completion Callbacks in C++

The `begin_` method is overloaded to allow you to provide completion callbacks. Here are the corresponding methods for the `getName` operation:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::AsyncResultPtr begin_getName(
                        Ice::Int number,
                        const Ice::CallbackPtr& __del,
                        const Ice::LocalObjectPtr& __cookie = 0);

Ice::AsyncResultPtr begin_getName(
                        Ice::Int number,
                        const Ice::Context& __ctx,
                        const Ice::CallbackPtr& __del,
                        const Ice::LocalObjectPtr& __cookie = 0);
```

The second version of `begin_getName` lets you override the default context. (We discuss the purpose of the `cookie` parameter in the next section.) Following the in-parameters, the `begin_` method accepts a parameter of type `Ice::CallbackPtr`. This is a smart pointer to a callback class that is provided by the Ice run time. This class stores an instance of a callback class that you implement. The Ice run time invokes a method on your callback instance when an asynchronous operation completes. Your callback class must provide a method that returns `void` and accepts a single parameter of type `const AsyncResultPtr&`, for example:

C++

```
class MyCallback : public IceUtil::Shared {
public:
    void finished(const Ice::AsyncResultPtr& r) {
        EmployeesPrx e = EmployeesPrx::uncheckedCast(r->getProxy());
        try {
            string name = e->end_getName(r);
            cout << "Name is: " << name << endl;
        } catch (const Ice::Exception& ex) {
            cerr << "Exception is: " << ex << endl;
        }
    }
};
typedef IceUtil::Handle<MyCallback> MyCallbackPtr;
```

Note that your callback class must derive from `IceUtil::Shared`. The callback method can have any name you prefer but its signature must match the preceding example.

The implementation of your callback method must call the `end_` method. The proxy for the call is available via the `getProxy` method on the `AsyncResult` that is passed by the Ice run time. The return type of `getProxy` is `Ice::ObjectPrx`, so you must down-cast the proxy to its correct type. (You should always use an `uncheckedCast` to do this, otherwise you will send an additional message to the server to verify the proxy type.)

Your callback method should catch and handle any exceptions that may be thrown by the `end_` method. If you allow an exception to escape from the callback method, the Ice run time produces a log entry by default and ignores the exception. (You can disable the log message by setting the property `Ice.Warn.AMICallback` to zero.)

To inform the Ice run time that you want to receive a callback for the completion of the asynchronous call, you pass the callback instance to the `begin_` method:

Copyright 2016, ZeroC, Inc.

C++

```
EmployeesPrx e = ...;

MyCallbackPtr cb = new MyCallback;
Ice::CallbackPtr d = Ice::newCallback(cb, &MyCallback::finished);

e->begin_getName(99, d);
```

Note the call to `Ice::newCallback` in this example. This helper function expects a smart pointer to your callback instance and a member function pointer that specifies your callback method.

Using Cookies for Generic Completion Callbacks in C++

It is common for the `end_` method to require access to some state that is established by the code that calls the `begin_` method. As an example, consider an application that asynchronously starts a number of operations and, as each operation completes, needs to update different user interface elements with the results. In this case, the `begin_` method knows which user interface element should receive the update, and the `end_` method needs access to that element.

The API allows you to pass such state by providing a cookie. A cookie is an instance of a class that you write; the class can contain whatever data you want to pass, as well as any methods you may want to add to manipulate that data.

The only requirement on the cookie class is that it must derive from `Ice::LocalObject`. Here is an example implementation that stores a `WidgetHandle`. (We assume that this class provides whatever methods are needed by the `end_` method to update the display.)

C++

```
class Cookie : public Ice::LocalObject
{
public:
    Cookie(WidgetHandle h) : _h(h) {}
    WidgetHandle getWidget() { return _h; }

private:
    WidgetHandle _h;
};
typedef IceUtil::Handle<Cookie> CookiePtr;
```

When you call the `begin_` method, you pass the appropriate cookie instance to inform the `end_` method how to update the display:

Copyright 2016, ZeroC, Inc.

C++

```
// Make cookie for call to getName(99).
CookiePtr cookie1 = new Cookie(widgetHandle1);

// Make cookie for call to getName(42);
CookiePtr cookie2 = new Cookie(widgetHandle2);

// Invoke the getName operation with different cookies.
e->begin_getName(99, getNameCB, cookie1);
e->begin_getName(24, getNameCB, cookie2);
```

The `end_` method can retrieve the cookie from the `AsyncResult` by calling `getCookie`. For this example, we assume that widgets have a `writeString` method that updates the relevant UI element:

C++

```
void
MyCallback::getName(const Ice::AsyncResultPtr& r)
{
    EmployeesPrx e = EmployeesPrx::uncheckedCast(r->getProxy());
    CookiePtr cookie = CookiePtr::dynamicCast(r->getCookie());
    try {
        string name = e->end_getName(r);
        cookie->getWidget()->writeString(name);
    } catch (const Ice::Exception& ex) {
        handleException(ex);
    }
}
```

The cookie provides a simple and effective way for you to pass state between the point where an operation is invoked and the point where its results are processed. Moreover, if you have a number of operations that share common state, you can pass the same cookie instance to multiple invocations.

Type-Safe Completion Callbacks in C++

The generic callback API is not entirely type-safe:

- You must down-cast the return value of `getProxy` to the correct proxy type before you can call the `end_` method.
- You must call the correct `end_` method to match the operation called by the `begin_` method.
- If you use a cookie, you must down-cast the cookie to the correct type before you can access the data inside the cookie.
- You must remember to catch exceptions when you call the `end_` method; if you forget to do this, you will not know that the operation failed.

`slice2cpp` generates an additional type-safe API that takes care of these chores for you. The type-safe API is provided as a template that works much like the `Ice::Callback` class of the generic API, but requires strongly-typed method signatures.

To use type-safe callbacks, you must implement a callback class that provides two callback methods:

- A success callback that is called if the operation succeeds
- A failure callback that is called if the operation raises an exception

As for the generic API, your callback class must derive from `IceUtil::Shared`. Here is a callback class for an invocation of the `getName` operation:

Copyright 2016, ZeroC, Inc.

C++

```
class MyCallback : public IceUtil::Shared
{
public:
    void getNameCB(const string& name) {
        cout << "Name is: " << name << endl;
    }

    void failureCB(const Ice::Exception& ex) {
        cerr << "Exception is: << ex << endl;
    }
};
```

The callback methods can have any name you prefer and must have `void` return type. The failure callback always has a single parameter of type `const Ice::Exception&`. The success callback parameters depend on the operation signature. If the operation has non-`void` retur n type, the first parameter of the success callback is the return value. The return value (if any) is followed by a parameter for each out-parameter of the corresponding Slice operation, in the order of declaration.

To receive these callbacks, you instantiate your callback instance and specify the methods you have defined before passing a smart pointer to a callback wrapper instance to the `begin_` method:

C++

```
MyCallbackPtr cb = new MyCallback;

Callback_Employees_getNamePtr getNameCB =
    newCallback_Employees_getName(cb, &MyCallback::getNameCB, &MyCallback::failureC
B);

Callback_Employees_getNumberPtr getNumberCB =
    newCallback_Employees_getNumber(cb, &MyCallback::getNumberCB, &MyCallback::fail
ureCB);

e->begin_getName(99, getNameCB);
e->begin_getNumber("Fred", getNumberCB);
```

Note how this code creates instances of two smart pointer types generated by `slice2cpp` named `Callback_Employees_getNamePtr` a nd `Callback_Employees_getNumberPtr`. Each smart pointer points to a template instance that encapsulates your callback instance and two member function pointers for the callback methods. The name of this smart pointer type is formed as follows:

<*module*>::Callback_<*interface*>_<*operation*>Ptr

Also note that the code uses helper functions to initialize the smart pointers. The first argument to the helper function is your callback instance, and the two following arguments are the success and failure member function pointers, respectively. The name of this helper function is formed as follows:

<*module*>::newCallback_<*interface*>_<*operation*>

It is legal to pass a null pointer as the success or failure callback. For the success callback, this is legal only for operations that have `void` re turn type and no out-parameters. This is useful if you do not care when the operation completes but want to know if the call failed. If you pass a null exception callback, the Ice run time will ignore any exception that is raised by the invocation.

The type of the success and exception member function pointers is provided as `Response` and `Exception` typedefs by the callback template. For example, you can ignore exceptions for an invocation of `getName` as follows:

Copyright 2016, ZeroC, Inc.

C++

```
Callback_Employees_getName::Exception nullException = 0;

MyCallbackPtr cb = new MyCallback;

Callback_Employees_getNamePtr getNameCB =
    newCallback_Employees_getName(cb, &MyCallback::getNameCB, nullException);

e->begin_getName(99, getNameCB); // Ignores exceptions
```

Using Cookies for Type-Safe Completion Callbacks in C++

The `begin_` method optionally accepts a cookie as a trailing parameter. As for the generic API, you can use the cookie to share state between the `begin_` and `end_` methods. However, with the type-safe API, there is no need to down-cast the cookie. Instead, the cookie parameter that is passed to the `end_` method is strongly typed. Assuming that you have defined a `Cookie` class and `CookiePtr` smart pointer, you can pass a cookie to the `begin_` method as follows:

C++

```
MyCallbackPtr cb = new MyCallback;

Callback_Employees_getNamePtr getNameCB =
    newCallback_Employees_getName(cb, &MyCallback::getNameCB, &MyCallback::failureC
B);

CookiePtr cookie = new Cookie(widgetHandle);
e->begin_getName(99, getNameCB, cookie);
```

The callback methods of your callback class simply add the cookie parameter:

C++

```
class MyCallback : public IceUtil::Shared
{
public:
    void getNameCB(const string& name, const CookiePtr& cookie) {
        cookie->getWidget()->writeString(name);
    }

    void failureCB(const Ice::Exception& ex, const CookiePtr& cookie) {
        cookie->getWidget()->writeError(ex.what());
    }
};
```

Lambda Completion Callbacks in C++

Copyright 2016, ZeroC, Inc.

`slice2cpp` generates overloaded versions of the `begin_` methods that accept lambda functions, eliminating the need to create a separate callback object. The Slice compiler generates the following methods for the `getName` operation:

C++

```
Ice::AsyncResultPtr
begin_getName(Ice::Int number,
              const Function<void (const std::string&)>& response,
              const Function<void (const Ice::Exception&)>& exception = ...,
              const Function<void (bool)>& sent = ...);

Ice::AsyncResultPtr
begin_getName(Ice::Int number, const Ice::Context& ctx,
              const Function<void (const std::string&)>& response,
              const Function<void (const Ice::Exception&)>& exception = ...,
              const Function<void (bool)>& sent = ...);

Ice::AsyncResultPtr
begin_getName(Ice::Int number,
              const Function<void (const Ice::AsyncResultPtr&)>& completed,
              const Function<void (const Ice::AsyncResultPtr&)>& sent = ...);

Ice::AsyncResultPtr
begin_getName(Ice::Int number, const Ice::Context& ctx,
              const Function<void (const Ice::AsyncResultPtr&)>& completed,
              const Function<void (const Ice::AsyncResultPtr&)>& sent = ...);
```

The overloaded versions of `begin_getName` allow you to use the type-safe and generic styles, with and without a context argument. Notice that default values are provided for the `exception` and `sent` arguments, meaning you can omit these arguments if you wish to ignore their events. The first two functions above correspond to the type-safe API and the second two have generic API semantics.

These `begin_` methods are generated by default, but can only be used with a compiler that supports C++11 lambda functions, and with an Ice installation that was built with C++11 features enabled.

We can invoke `begin_getName` as shown below:

C++

```
EmployeesPrx e = ...;

// Type-safe
e->begin_getName(42, [](const std::string& name) { cout << "name = " << name <<
endl; });

// Generic
e->begin_getName(42, [=](const Ice::AsyncResultPtr& r) {
    cout << "name = " << e->end_getName(r) << endl;
});
```

For the sake of brevity we have omitted the error handling that would normally be necessary.

Copyright 2016, ZeroC, Inc.

Asynchronous Oneway Invocations in C++

You can invoke operations via oneway proxies asynchronously, provided the operation has `void` return type, does not have any out-parameters, and does not raise user exceptions. If you call the `begin_` method on a oneway proxy for an operation that returns values or raises a user exception, the `begin_` method throws an `IceUtil::IllegalArgumentException`.

For the generic API, the callback method looks exactly as for a twoway invocation. However, for oneway invocations, the Ice run time does not call the callback method unless the invocation raised an exception during the `begin_` method ("on the way out").

For the type-safe API, the `newCallback` helper for `void` operations is overloaded so you can omit the success callback. For example, here is how you could call `ice_ping` asynchronously:

C++

```
MyCallbackPtr cb = new MyCallback;

Ice::Callback_Object_ice_pingPtr callback =
    Ice::newCallback_Object_ice_ping(cb, &MyCallback::failureCB);

p->begin_opVoid(callback);
```

Flow Control in C++

Asynchronous method invocations never block the thread that calls the `begin_` method: the Ice run time checks to see whether it can write the request to the local transport. If it can, it does so immediately in the caller's thread. (In that case, `AsyncResult::sentSynchronously` returns true.) Alternatively, if the local transport does not have sufficient buffer space to accept the request, the Ice run time queues the request internally for later transmission in the background. (In that case, `AsyncResult::sentSynchronously` returns false.)

This creates a potential problem: if a client sends many asynchronous requests at the time the server is too busy to keep up with them, the requests pile up in the client-side run time until, eventually, the client runs out of memory.

The API provides a way for you to implement flow control by counting the number of requests that are queued so, if that number exceeds some threshold, the client stops invoking more operations until some of the queued operations have drained out of the local transport.

For the generic API, you can create an additional callback method:

C++

```
class MyCallback : public IceUtil::Shared {
public:
    void finished(const Ice::AsyncResultPtr&);
    void sent(const Ice::AsyncResultPtr&);
};
typedef IceUtil::Handle<MyCallback> MyCallbackPtr;
```

As with any other callback method, you are free to choose any name you like. For this example, the name of the callback method is `sent`. You inform the Ice run time that you want to be informed when a call has been passed to the local transport by specifying the `sent` method as an additional parameter when you create the `Ice::Callback`:

Copyright 2016, ZeroC, Inc.

C++

```
EmployeesPrx e = ...;

MyCallbackPtr cb = new MyCallback;
Ice::CallbackPtr d = Ice::newCallback(cb, &MyCallback::finished,
&MyCallback::sent);

e->begin_getName(99, d);
```

If the Ice run time can immediately pass the request to the local transport, it does so and invokes the `sent` method from the thread that calls the `begin_` method. On the other hand, if the run time has to queue the request, it calls the `sent` method from a different thread once it has written the request to the local transport. In addition, you can find out from the `AsyncResult` that is returned by the `begin_` method whether the request was sent synchronously or was queued, by calling `sentSynchronously`.

For the generic API, the `sent` method has the following signature:

C++

```
void sent(const Ice::AsyncResult&);
```

For the type-safe API, there are two versions, one without and one with a cookie:

C++

```
void sent(bool sentSynchronously);
void sent(bool sentSynchronously, const <CookiePtr>& cookie);
```

For the version with a cookie, *<CookiePtr>* is replaced with the actual type of the cookie smart pointer you passed to the `begin_` method.

The `sent` methods allow you to limit the number of queued requests by counting the number of requests that are queued and decrementing the count when the Ice run time passes a request to the local transport.

Asynchronous Batch Requests in C++

Applications that send batched requests can either flush a batch explicitly or allow the Ice run time to flush automatically. The proxy method `ice_flushBatchRequests` performs an immediate flush using the synchronous invocation model and may block the calling thread until the entire message can be sent. Ice also provides asynchronous versions of this method so you can flush batch requests asynchronously.

`begin_ice_flushBatchRequests` and `end_ice_flushBatchRequests` are proxy methods that flush any batch requests queued by that proxy.

In addition, similar methods are available on the communicator and the `Connection` object that is returned by `AsyncResult::getConnection`. These methods flush batch requests sent via the same communicator and via the same connection, respectively.

Concurrency Semantics for AMI in C++

The Ice run time always invokes your callback methods from a separate thread. This means that you can safely use a non-recursive mutex without risking deadlock. There is one exception to this rule: the run time calls the `sent` callback from the thread calling the `begin_` method if the request could be sent synchronously. In the `sent` callback, you know which thread is calling the callback by looking at the `sentSynchronously` member or parameter, so you can take appropriate action to avoid a self-deadlock.

Copyright 2016, ZeroC, Inc.

See Also

- C++ Mapping for Classes
- Smart Pointers for Classes
- Request Contexts
- Batched Invocations
- Collocated Invocation and Dispatch

Copyright 2016, ZeroC, Inc.

Using Slice Checksums in C++

The Slice compilers can optionally generate checksums of Slice definitions. For `slice2cpp`, the `--checksum` option causes the compiler to generate code in each C++ source file that accumulates checksums in a global map. A copy of this map can be obtained by calling a function defined in the header file `Ice/SliceChecksums.h`:

C++

```
namespace Ice {
    Ice::SliceChecksumDict sliceChecksums();
}
```

In order to verify a server's checksums, a client could simply compare the dictionaries using the equality operator. However, this is not feasible if it is possible that the server might be linked with more Slice definitions than the client. A more general solution is to iterate over the local checksums as demonstrated below:

C++

```
Ice::SliceChecksumDict serverChecksums = ...
Ice::SliceChecksumDict localChecksums = Ice::sliceChecksums();

for (Ice::SliceChecksumDict::const_iterator p = localChecksums.begin();
     p != localChecksums.end(); ++p) {

    Ice::SliceChecksumDict::const_iterator q = serverChecksums.find(p->first);
    if (q == serverChecksums.end()) {
        // No match found for type id!
    } else if (p->second != q->second) {
        // Checksum mismatch!
    }
}
```

In this example, the client first verifies that the server's dictionary contains an entry for each Slice type ID, and then it proceeds to compare the checksums.

See Also

- Slice Checksums
- slice2cpp Command-Line Options

Copyright 2016, ZeroC, Inc.

Example of a File System Client in C++

This page presents a very simple client to access a server that implements the file system we developed in Slice for a Simple File System. The C++ code shown here hardly differs from the code you would write for an ordinary C++ program. This is one of the biggest advantages of using Ice: accessing a remote object is as easy as accessing an ordinary, local C++ object. This allows you to put your effort where you should, namely, into developing your application logic instead of having to struggle with arcane networking APIs. This is true for the server side as well, meaning that you can develop distributed applications easily and efficiently.

We now have seen enough of the client-side C++ mapping to develop a complete client to access our remote file system. For reference, here is the Slice definition once more:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };

    exception GenericError {
        string reason;
    };

    sequence<string> Lines;

    interface File extends Node {
        idempotent Lines read();
        idempotent void write(Lines text) throws GenericError;
    };

    sequence<Node*> NodeSeq;

    interface Directory extends Node {
        idempotent NodeSeq list();
    };
};
```

To exercise the file system, the client does a recursive listing of the file system, starting at the root directory. For each node in the file system, the client shows the name of the node and whether that node is a file or directory. If the node is a file, the client retrieves the contents of the file and prints them.

The body of the client code looks as follows:

C++

Copyright 2016, ZeroC, Inc.

```
#include <Ice/Ice.h>
#include <Filesystem.h>
#include <iostream>
#include <iterator>

using namespace std;
using namespace Filesystem;

static void
listRecursive(const DirectoryPrx& dir, int depth = 0)
{
    // ...
}

int
main(int argc, char* argv[])
{
    int status = 0;
    Ice::CommunicatorPtr ic;
    try {
        // Create a communicator
        //
        ic = Ice::initialize(argc, argv);

        // Create a proxy for the root directory
        //
        Ice::ObjectPrx base = ic->stringToProxy("RootDir:default -p 10000");
        if (!base)
            throw "Could not create proxy";

        // Down-cast the proxy to a Directory proxy
        //
        DirectoryPrx rootDir = DirectoryPrx::checkedCast(base);
        if (!rootDir)
            throw "Invalid proxy";

        // Recursively list the contents of the root directory
        //
        cout << "Contents of root directory:" << endl;
        listRecursive(rootDir);
    } catch (const Ice::Exception& ex) {
        cerr << ex << endl;
        status = 1;
    } catch (const char* msg) {
        cerr << msg << endl;
        status = 1;
    }

    // Clean up
    //
    if (ic)
        ic->destroy();

    return status;
}
```


Copyright 2016, ZeroC, Inc.

1. The code includes a few header files:
   - `Ice/Ice.h`:
     Always included in both client and server source files, provides definitions that are necessary for accessing the Ice run time.
   - `Filesystem.h`:
     The header that is generated by the Slice compiler from the Slice definitions in `Filesystem.ice`.
   - `iostream`:
     The client uses the `iostream` library to produce its output.
   - `iterator`:
     The implementation of `listRecursive` uses an STL iterator.
2. The code adds `using` declarations for the `std` and `Filesystem` namespaces.
3. The structure of the code in `main` follows what we saw in Hello World Application. After initializing the run time, the client creates a proxy to the root directory of the file system. For this example, we assume that the server runs on the local host and listens using the default protocol (TCP/IP) at port 10000. The object identity of the root directory is known to be `RootDir`.
4. The client down-casts the proxy to `DirectoryPrx` and passes that proxy to `listRecursive`, which prints the contents of the file system.

Most of the work happens in `listRecursive`:

C++

```
// Recursively print the contents of directory "dir" in
// tree fashion. For files, show the contents of each file.
// The "depth" parameter is the current nesting level
// (for indentation).

static void
listRecursive(const DirectoryPrx& dir, int depth = 0)
{
    string indent(++depth, '\t');

    NodeSeq contents = dir->list();

    for (NodeSeq::const_iterator i = contents.begin(); i != contents.end(); ++i) {
        DirectoryPrx dir = DirectoryPrx::checkedCast(*i);
        FilePrx file = FilePrx::uncheckedCast(*i);
        cout << indent << (*i)->name()
<< (dir ? " (directory):" : " (file):") << endl;
        if (dir) {
            listRecursive(dir, depth);
        } else {
            Lines text = file->read();
            for (Lines::const_iterator j = text.begin(); j != text.end(); ++j) {
                cout << indent << "\t" << *j << endl;
            }
        }
    }
}
```

The function is passed a proxy to a directory to list, and an indent level. (The indent level increments with each recursive call and allows the code to print the name of each node at an indent level that corresponds to the depth of the tree at that node.) `listRecursive` calls the list operation on the directory and iterates over the returned sequence of nodes:

1. The code does a `checkedCast` to narrow the `Node` proxy to a `Directory` proxy, as well as an `uncheckedCast` to narrow the `No de` proxy to a `File` proxy. Exactly one of those casts will succeed, so there is no need to call `checkedCast` twice: if the `Node` *is-a* `Directory`, the code uses the `DirectoryPrx` returned by the `checkedCast`; if the `checkedCast` fails, we *know* that the `Node` *i*

Copyright 2016, ZeroC, Inc.

s-a `File` and, therefore, an `uncheckedCast` is sufficient to get a `FilePrx`.
In general, if you know that a down-cast to a specific type will succeed, it is preferable to use an `uncheckedCast` instead of a `checkedCast` because an `uncheckedCast` does not incur any network traffic.

2. The code prints the name of the file or directory and then, depending on which cast succeeded, prints `"(directory)"` or `"(file)"` following the name.
3. The code checks the type of the node:
    - If it is a directory, the code recurses, incrementing the indent level.
    - If it is a file, the code calls the `read` operation on the file to retrieve the file contents and then iterates over the returned sequence of lines, printing each line.

Assume that we have a small file system consisting of two files and a directory as follows:

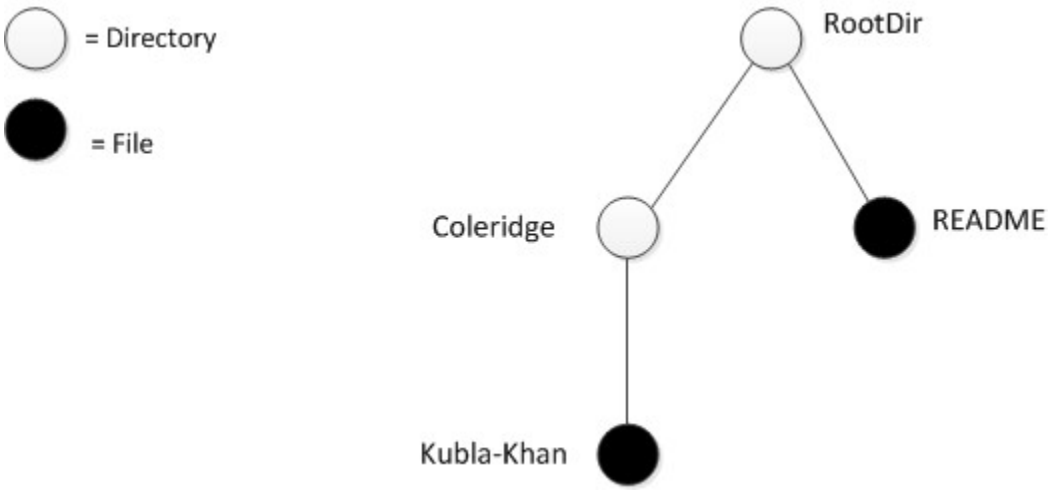

A small file system.

The output produced by the client for this file system is:

```
Contents of root directory:
        README (file):
                This file system contains a collection of poetry.
        Coleridge (directory):
                Kubla_Khan (file):
                        In Xanadu did Kubla Khan
                        A stately pleasure-dome decree:
                        Where Alph, the sacred river, ran
                        Through caverns measureless to man
                        Down to a sunless sea.
```

Note that, so far, our client (and server) are not very sophisticated:

- The protocol and address information are hard-wired into the code.
- The client makes more remote procedure calls than strictly necessary; with minor redesign of the Slice definitions, many of these calls can be avoided.

We will see how to address these shortcomings in our discussions of IceGrid and object life cycle.

See Also

- Hello World Application
- Slice for a Simple File System
- Example of a File System Server in C++
- Object Life Cycle
- IceGrid

Copyright 2016, ZeroC, Inc.

Server-Side Slice-to-C++ Mapping

The mapping for Slice data types to C++ is identical on the client side and server side. This means that everything in Client-Side Slice-to-C++ Mapping also applies to the server side. However, for the server side, there are a few additional things you need to know — specifically how to:

- Initialize and finalize the server-side run time
- Implement servants
- Pass parameters and throw exceptions
- Create servants and register them with the Ice run time.

Because the mapping for Slice data types is identical for clients and servers, the server-side mapping only adds a few additional mechanisms to the client side: a small API to initialize and finalize the run time, plus a few rules for how to derive servant classes from skeletons and how to register servants with the server-side run time.

Although the examples we present are very simple, they accurately reflect the basics of writing an Ice server. Of course, for more sophisticated servers, you will be using additional APIs, for example, to improve performance or scalability. However, these APIs are all described in Slice, so, to use these APIs, you need not learn any C++ mapping rules beyond those we describe here.

Topics

- The Server-Side main Function in C++
- Server-Side C++ Mapping for Interfaces
- Parameter Passing in C++
- Raising Exceptions in C++
- Object Incarnation in C++
- Asynchronous Method Dispatch (AMD) in C++
- Example of a File System Server in C++

Copyright 2016, ZeroC, Inc.

The Server-Side main Function in C++

On this page:

- A Basic main Function in C++
- The Ice::Application Class
  - Using Ice::Application on the Client Side
  - Catching Signals in C++
  - Ice::Application and Properties
  - Limitations of Ice::Application
- The Ice::Service Class
  - Ice::Service Member Functions
  - Unix Daemons
  - Windows Services
  - Ice::Service Logging Considerations

A Basic `main` Function in C++

The main entry point to the Ice run time is represented by the local Slice interface `Ice::Communicator`. As for the client side, you must initialize the Ice run time by calling `Ice::initialize` before you can do anything else in your server. `Ice::initialize` returns a smart pointer to an instance of an `Ice::Communicator`:

C++

```
int
main(int argc, char* argv[])
{
    Ice::CommunicatorPtr ic = Ice::initialize(argc, argv);
    // ...
}
```

`Ice::initialize` accepts a C++ reference to `argc` and `argv`. The function scans the argument vector for any command-line options that are relevant to the Ice run time; any such options are removed from the argument vector so, when `Ice::initialize` returns, the only options and arguments remaining are those that concern your application. If anything goes wrong during initialization, `initialize` throws an exception.

`Ice::initialize` has additional overloads to permit other information to be passed to the Ice run time.

Before leaving your `main` function, you must call `Communicator::destroy`. The `destroy` operation is responsible for finalizing the Ice run time. In particular, `destroy` waits for any operation implementations that are still executing in the server to complete. In addition, `destroy` ensures that any outstanding threads are joined with and reclaims a number of operating system resources, such as file descriptors and memory. Never allow your `main` function to terminate without calling `destroy` first; doing so has undefined behavior.

The general shape of our server-side `main` function is therefore as follows:

Copyright 2016, ZeroC, Inc.

C++

```
#include <Ice/Ice.h>

int
main(int argc, char* argv[])
{
    int status = 0;
    Ice::CommunicatorPtr ic;
    try {
        ic = Ice::initialize(argc, argv);

        // Server code here...

    } catch (const Ice::Exception& e) {
        cerr << e << endl;
        status = 1;
    } catch (const std::string& msg) {
        cerr << msg << endl;
        status = 1;
    } catch (const char* msg) {
        cerr << msg << endl;
        status = 1;
    }
    if (ic) {
        try {
            ic->destroy();
        } catch (const Ice::Exception& e) {
            cerr << e << endl;
            status = 1;
        }
    }
    return status;
}
```

Note that the code places the call to `Ice::initialize` in to a `try` block and takes care to return the correct exit status to the operating system. Also note that an attempt to destroy the communicator is made only if the initialization succeeded.

The `catch` handlers for `const std::string &` and `const char *` are in place as a convenience feature: if we encounter a fatal error condition anywhere in the server code, we can simply throw a string or a string literal containing an error message; this causes the stack to be unwound back to `main`, at which point the error message is printed and, after destroying the communicator, `main` terminates with non-zero exit status.

The `Ice::Application` Class

The preceding structure for the `main` function is so common that Ice offers a class, `Ice::Application`, that encapsulates all the correct initialization and finalization activities. The definition of the class is as follows (with some detail omitted for now):

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {
    enum SignalPolicy { HandleSignals, NoSignalHandling };

    class Application /* ... */ {
    public:
        Application(SignalPolicy = HandleSignals);
        virtual ~Application();

        int main(int argc, char*[] argv);
        int main(int argc, char*[] argv, const char* config);
        int main(int argc, char*[] argv, const Ice::InitializationData& id);
        int main(int argc, char* const [] argv);
        int main(int argc, char* const [] argv, const char* config);
        int main(int argc, char* const [] argv, const Ice::InitializationData& id);
        int main(const Ice::StringSeq&);
        int main(const Ice::StringSeq&, const char* config);
        int main(const Ice::StringSeq&, const Ice::InitializationData& id);

#ifdef _WIN32
        int main(int argc, wchar_t*[] argv);
        int main(int argc, wchar_t*[] argv, const char* config);
        int main(int argc, wchar_t*[] argv, const Ice::InitializationData& id);
#endif

        virtual int run(int, char*[]) = 0;

        static  const char* appName();
        static  CommunicatorPtr communicator();
        // ...
    };
}
```

The intent of this class is that you specialize `Ice::Application` and implement the pure virtual `run` method in your derived class. Whatever code you would normally place in `main` goes into the `run` method instead. Using `Ice::Application`, our program looks as follows:

Copyright 2016, ZeroC, Inc.

C++

```
#include <Ice/Ice.h>

class MyApplication : virtual public Ice::Application {
public:
    virtual int run(int, char*[]) {

        // Server code here...

        return 0;
    }
};

int
main(int argc, char* argv[])
{
    MyApplication app;
    return app.main(argc, argv);
}
```

Note that `Application::main` is overloaded: you can pass a string sequence instead of an `argc`/`argv` pair. This is useful if you need to parse application-specific property settings on the command line. You also can call `main` with an optional file name or an `InitializationData` structure.

If you pass a configuration file name to `main`, the property settings in this file are overridden by settings in a file identified by the `ICE_CONFIG` environment variable (if defined). Property settings supplied on the command line take precedence over all other settings.

The `Application::main` function does the following:

1. It installs an exception handler for `Ice::Exception`. If your code fails to handle an Ice exception, `Application::main` prints the exception details on `stderr` before returning with a non-zero return value.
2. It installs exception handlers for `const std::string &` and `const char*`. This allows you to terminate your server in response to a fatal error condition by throwing a `std::string` or a string literal. `Application::main` prints the string on `stderr` before returning a non-zero return value.
3. It initializes (by calling `Ice::initialize`) and finalizes (by calling `Communicator::destroy`) a communicator. You can get access to the communicator for your server by calling the static `communicator()` member function.
4. It scans the argument vector for options that are relevant to the Ice run time and removes any such options. The argument vector that is passed to your `run` method therefore is free of Ice-related options and only contains options and arguments that are specific to your application.
5. It provides the name of your application via the static `appName` member function. The return value from this call is `argv[0]`, so you can get at `argv[0]` from anywhere in your code by calling `Ice::Application::appName` (which is often necessary for error messages).
6. It installs a signal handler that properly destroys the communicator.
7. It installs a per-process logger if the application has not already configured one. The per-process logger uses the value of the `Ice.ProgramName` property as a prefix for its messages and sends its output to the standard error channel. An application can also specify an alternate logger.

Using `Ice::Application` ensures that your program properly finalizes the Ice run time, whether your server terminates normally or in response to an exception or signal. We recommend that all your programs use this class; doing so makes your life easier. In addition, `Ice::Application` also provides features for signal handling and configuration that you do not have to implement yourself when you use this class.

Using `Ice::Application` on the Client Side

You can use `Ice::Application` for your clients as well: simply implement a class that derives from `Ice::Application` and place the client code into its `run` method. The advantage of this approach is the same as for the server side: `Ice::Application` ensures that the communicator is destroyed correctly even in the presence of exceptions.

Copyright 2016, ZeroC, Inc.

Catching Signals in C++

The simple server we developed in Hello World Application had no way to shut down cleanly: we simply interrupted the server from the command line to force it to exit. Terminating a server in this fashion is unacceptable for many real-life server applications: typically, the server has to perform some cleanup work before terminating, such as flushing database buffers or closing network connections. This is particularly important on receipt of a signal or keyboard interrupt to prevent possible corruption of database files or other persistent data.

To make it easier to deal with signals, `Ice::Application` encapsulates the platform-independent signal handling capabilities provided by the class `IceUtil::CtrlCHandler`. This allows you to cleanly shut down on receipt of a signal and to use the same source code regardless of the underlying operating system and threading package:

C++

```
namespace Ice {
    class Application : /* ... */ {
    public:
        // ...
        static void destroyOnInterrupt();
        static void shutdownOnInterrupt();
        static void ignoreInterrupt();
        static void callbackOnInterrupt();
        static void holdInterrupt();
        static void releaseInterrupt();
        static bool interrupted();

        virtual void interruptCallback(int);
    };
}
```

You can use `Ice::Application` under both Windows and Unix: for Unix, the member functions control the behavior of your application for `SIGINT`, `SIGHUP`, and `SIGTERM`; for Windows, the member functions control the behavior of your application for `CTRL_C_EVENT`, `CTRL_BREAK_EVENT`, `CTRL_CLOSE_EVENT`, `CTRL_LOGOFF_EVENT`, and `CTRL_SHUTDOWN_EVENT`.

The functions behave as follows:

- `destroyOnInterrupt`
  This function creates an `IceUtil::CtrlCHandler` that destroys the communicator when one of the monitored signals is raised. This is the default behavior.
- `shutdownOnInterrupt`
  This function creates an `IceUtil::CtrlCHandler` that shuts down the communicator when one of the monitored signals is raised.
- `ignoreInterrupt`
  This function causes signals to be ignored.
- `callbackOnInterrupt`
  This function configures `Ice::Application` to invoke `interruptCallback` when a signal occurs, thereby giving the subclass responsibility for handling the signal. Note that if the signal handler needs to terminate the program, you must call `_exit` (instead of `exit`). This prevents global destructors from running which, depending on the activities of other threads in the program, could cause deadlock or assertion failures.
- `holdInterrupt`
  This function causes signals to be held.
- `releaseInterrupt`
  This function restores signal delivery to the previous disposition. Any signal that arrives after `holdInterrupt` was called is delivered when you call `releaseInterrupt`.
- `interrupted`
  This function returns `true` if a signal caused the communicator to shut down, `false` otherwise. This allows us to distinguish intentional shutdown from a forced shutdown that was caused by a signal. This is useful, for example, for logging purposes.

Copyright 2016, ZeroC, Inc.

- `interruptCallback`
  A subclass overrides this function to respond to signals. The Ice run time may call this function concurrently with any other thread. If the function raises an exception, the Ice run time prints a warning on `cerr` and ignores the exception.

By default, `Ice::Application` behaves as if `destroyOnInterrupt` was invoked, therefore our server `main` function requires no change to ensure that the program terminates cleanly on receipt of a signal. (You can disable the signal-handling functionality of `Ice::Application` by passing the enumerator `NoSignalHandling` to the constructor. In that case, signals retain their default behavior, that is, terminate the process.) However, we add a diagnostic to report the occurrence of a signal, so our `main` function now looks like:

C++

```
#include <Ice/Ice.h>

class MyApplication : virtual public Ice::Application {
public:
    virtual int run(int, char*[]) {

        // Server code here...

        if (interrupted())
            cerr << appName() << ": terminating" << endl;

        return 0;
    }
};

int
main(int argc, char* argv[])
{
    MyApplication app;
    return app.main(argc, argv);
}
```

Note that, if your server is interrupted by a signal, the Ice run time waits for all currently executing operations to finish. This means that an operation that updates persistent state cannot be interrupted in the middle of what it was doing and cause partial update problems.

Under Unix, if you handle signals with your own handler (by deriving a subclass from `Ice::Application` and calling `callbackOnInterrupt`), the handler is invoked synchronously from a separate thread. This means that the handler can safely call into the Ice run time or make system calls that are not async-signal-safe without fear of deadlock or data corruption. Note that `Ice::Application` blocks delivery of `SIGINT`, `SIGHUP`, and `SIGTERM`. If your application calls `exec`, this means that the child process will also ignore these signals; if you need the default behavior of these signals in the `exec`'d process, you must explicitly reset them to `SIG_DFL` before calling `exec`.

`Ice::Application` and Properties

Apart from the functionality shown in this section, `Ice::Application` also takes care of initializing the Ice run time with property values. Properties allow you to configure the run time in various ways. For example, you can use properties to control things such as the thread pool size or port number for a server.

Limitations of `Ice::Application`

`Ice::Application` is a singleton class that creates a single communicator. If you are using multiple communicators, you cannot use `Ice::Application`. Instead, you must structure your code as we saw in Hello World Application (taking care to always destroy the communicators).

The `Ice::Service` Class

The `Ice::Application` class is very convenient for general use by Ice client and server applications. In some cases, however, an application may need to run at the system level as a Unix daemon or Windows service. For these situations, Ice includes `Ice::Service`, a

Copyright 2016, ZeroC, Inc.

singleton class that is comparable to `Ice::Application` but also encapsulates the low-level, platform-specific initialization and shutdown procedures common to system services. The `Ice::Service` class is defined as follows:

C++

Copyright 2016, ZeroC, Inc.

```
namespace Ice {
    class Service {
    public:
        Service();

        virtual bool shutdown();
        virtual void interrupt();

        int main(int& argc, char* argv[],
                 const Ice::InitializationData& = Ice::InitializationData());
        int main(Ice::StringSeq& args,
                 const Ice::InitializationData& = Ice::InitializationData());

        Ice::CommunicatorPtr communicator() const;

        static Service* instance();

        bool service() const;
        std::string name() const;
        bool checkSystem() const;

        int run(int& argc, char* argv[], const Ice::InitializationData&);

#ifdef _WIN32
        int main(int& argc, wchar_t* argv[], const InitializationData& =
InitializationData());

        void configureService(const std::string& name);
#else
        void configureDaemon(bool changeDir, bool closeFiles,
const std::string& pidFile);
#endif

        virtual void handleInterrupt(int);

    protected:
        virtual bool start(int argc, char* argv[], int& status) = 0;
        virtual void waitForShutdown();
        virtual bool stop();
        virtual Ice::CommunicatorPtr initializeCommunicator(
            int& argc, char* argv[],
            const Ice::InitializationData&);

        virtual void syserror(const std::string& msg);
        virtual void error(const std::string& msg);
        virtual void warning(const std::string& msg);
        virtual void trace(const std::string& msg);
        virtual void print(const std::string& msg);

        void enableInterrupt();
        void disableInterrupt();

        // ...
    };
}
```


Copyright 2016, ZeroC, Inc.

At a minimum, an Ice application that uses the `Ice::Service` class must define a subclass and override the `start` member function, which is where the service must perform its startup activities, such as processing command-line arguments, creating an object adapter, and registering servants. The application's `main` function must instantiate the subclass and typically invokes its `main` member function, passing the program's argument vector as parameters. The example below illustrates a minimal `Ice::Service` subclass:

C++

```
#include <Ice/Service.h>

class MyService : public Ice::Service {
protected:
    virtual bool start(int, char*[], int&);
private:
    Ice::ObjectAdapterPtr _adapter;
};

bool
MyService::start(int argc, char* argv[], int& status)
{
    _adapter = communicator()->createObjectAdapter("MyAdapter");
    _adapter->addWithUUID(new MyServantI);
    _adapter->activate();
    status = EXIT_SUCCESS;
    return true;
}

int
main(int argc, char* argv[])
{
    MyService svc;
    return svc.main(argc, argv);
}
```

The `Service::main` member function performs the following sequence of tasks:

1. Scans the argument vector for reserved options that indicate whether the program should run as a system service and removes these options from the argument vector (`argc` is adjusted accordingly). Additional reserved options are supported for administrative tasks.
2. Configures the program for running as a system service (if necessary) by invoking `configureService` or `configureDaemon`, as appropriate for the platform.
3. Invokes the `run` member function and returns its result.

Note that, as for `Application::main`, `Service::main` is overloaded to accept a string sequence instead of an `argc`/`argv` pair. This is useful if you need to parse application-specific property settings on the command line.

For maximum portability, we strongly recommend that all initialization tasks be performed in the `start` member function and not in the global `main` function. For example, allocating resources in `main` can cause program instability for Unix daemons.

The `Service::run` member function executes the service in the steps shown below:

1. Installs a signal handler.
2. Invokes the `initializeCommunicator` member function to obtain a communicator. The communicator instance can be accessed using the `communicator` member function.
3. Invokes the `start` member function. If `start` returns `false` to indicate failure, `run` destroys the communicator and returns immediately using the exit status provided in `status`.
4. Invokes the `waitForShutdown` member function, which should block until `shutdown` is invoked.

Copyright 2016, ZeroC, Inc.

5. Invokes the `stop` member function. If `stop` returns `true`, `run` considers the application to have terminated successfully.
6. Destroys the communicator.
7. Gracefully terminates the system service (if necessary).

If an unhandled exception is caught by `Service::run`, a descriptive message is logged, the communicator is destroyed and the service is terminated.

`Ice::Service` Member Functions

The virtual member functions in `Ice::Service` represent the points at which a subclass can intercept the service activities. All of the virtual member functions (except `start`) have default implementations.

- `void handleInterrupt(int sig)`
  Invoked by the `CtrlCHandler` when a signal occurs. The default implementation ignores the signal if it represents a logoff event and the `Ice.Nohup` property is set to a value larger than zero, otherwise it invokes the `interrupt` member function.

- `Ice::CommunicatorPtr initializeCommunicator(int & argc, char * argv[],`
  `const Ice::InitializationData & data)`
  Initializes a communicator. The default implementation invokes `Ice::initialize` and passes the given arguments.

- `void interrupt()`
  Invoked by the signal handler to indicate a signal was received. The default implementation invokes the `shutdown` member function.

- `bool shutdown()`
  Causes the service to begin the shutdown process. The default implementation invokes `shutdown` on the communicator. The subclass must return `true` if shutdown was started successfully, and `false` otherwise.

- `bool start(int argc, char * argv[], int & status)`
  Allows the subclass to perform its startup activities, such as scanning the provided argument vector for recognized command-line options, creating an object adapter, and registering servants. The subclass must return `true` if startup was successful, and `false` o therwise. The subclass can set an exit status via the `status` parameter. This status is returned by `main`.

- `bool stop()`
  Allows the subclass to clean up prior to termination. The default implementation does nothing but return `true`. The subclass must return `true` if the service has stopped successfully, and `false` otherwise.

- `void syserror(const std::string & msg)`
  `void error(const std::string & msg)`
  `void warning(const std::string & msg)`
  `void trace(const std::string & msg)`
  `void print(const std::string & msg)`
  Convenience functions for logging messages to the communicator's logger. The `syserror` member function includes a description of the system's current error code. You can also log messages to these functions using utility classes similar to the C++ Logger Utility Classes: these classes are `ServiceSysError`, `ServiceError`, `ServiceWarning`, `ServiceTrace` and `ServicePrint`, all nested in the Service class.

- `void waitForShutdown()`
  Waits indefinitely for the service to shut down. The default implementation invokes `waitForShutdown` on the communicator.

The non-virtual member functions shown in the class definition are described below:

- `bool checkSystem() const`
  Returns true if the operating system supports Windows services or Unix daemons.

- `Ice::CommunicatorPtr communicator() const`
  Returns the communicator used by the service, as created by `initializeCommunicator`.

- `void configureDaemon(bool chdir, bool close, const std::string & pidFile)`
  Configures the program to run as a Unix daemon. The `chdir` parameter determines whether the daemon changes its working directory to the root directory. The `close` parameter determines whether the daemon closes unnecessary file descriptors (i.e., stdin, stdout, etc.). If a non-empty string is provided in the `pidFile` parameter, the daemon writes its process ID to the given file.

- `void configureService(const std::string & name)`
  Configures the program to run as a Windows service with the given name.

- `void disableInterrupt()`
  Disables the signal handling behavior in `Ice::Service`. When disabled, signals are ignored.

Copyright 2016, ZeroC, Inc.

- `void enableInterrupt()`
  Enables the signal handling behavior in `Ice::Service`. When enabled, the occurrence of a signal causes the `handleInterrupt` member function to be invoked.
- `static Service * instance()`
  Returns the singleton `Ice::Service` instance.
- ```
  int main(int & argc, char * argv[],
     const Ice::InitializationData & data = Ice::InitializationData())
  int main(Ice::StringSeq& args,
     const Ice::InitializationData& = Ice::InitializationData());
  int main(int & argc, wchar_t * argv[],
     const Ice::InitializationData & data = Ice::InitializationData())
  ```
  The primary entry point of the `Ice::Service` class. The tasks performed by this function are described earlier in this section. The function returns `EXIT_SUCCESS` for success, `EXIT_FAILURE` for failure. For Windows, this function is overloaded to allow you to pass a `wchar_t` argument vector.
- `std::string name() const`
  Returns the name of the service. If the program is running as a Windows service, the return value is the Windows service name, otherwise it returns the value of `argv[0]`.
- `int run(int & argc, char * argv[], const Ice::InitializationData & data)`
  Alternative entry point for applications that prefer a different style of service configuration. The program must invoke `configureService` (Windows) or `configureDaemon` (Unix) in order to run as a service. The tasks performed by this function were described earlier. The function normally returns `EXIT_SUCCESS` or `EXIT_FAILURE`, but the `start` method can also supply a different value via its `status` argument.
- `bool service() const`
  Returns true if the program is running as a Windows service or Unix daemon, or false otherwise.

Unix Daemons

On Unix platforms, passing `--daemon` causes your program to run as a daemon. When this option is present, `Ice::Service` performs the following additional actions:

- Creates a background child process in which `Service::main` performs its tasks. The foreground process does not terminate until the child process has successfully invoked the `start` member function. This behavior avoids the uncertainty often associated with starting a daemon from a shell script by ensuring that the command invocation does not complete until the daemon is ready to receive requests.
- Changes the current working directory of the child process to the root directory, unless `--nochdir` is specified.
- Closes all file descriptors, unless `--noclose` is specified. The standard input (stdin) channel is closed and reopened to `/dev/null`. Likewise, the standard output (stdout) and standard error (stderr) channels are also closed and reopened to `/dev/null` unless `Ice.StdOut` or `Ice.StdErr` are defined, respectively, in which case those channels use the designated log files.

  > The file descriptors are not closed until after the communicator is initialized, meaning standard input, standard output, and standard error are available for use during this time. For example, the IceSSL plug-in may need to prompt for a passphrase on standard input, or Ice may print the child's process id on standard output if the property `Ice.PrintProcessId` is set.

The following additional command-line options can be specified in conjunction with the `--daemon` option:

- `--pidfile` *`FILE`*
  This option writes the process ID of the service into the specified *`FILE`*.
- `--noclose`
  Prevents `Ice::Service` from closing unnecessary file descriptors. This can be useful during debugging and diagnosis because it provides access to the output from the daemon's standard output and standard error.
- `--nochdir`
  Prevents `Ice::Service` from changing the current working directory.

All of these options are removed from the argument vector that is passed to the `start` member function.

Copyright 2016, ZeroC, Inc.
  ```

We strongly recommend that you perform all initialization tasks in your service's `start` member function, and not in the global `main` function. This is especially important for process-specific resources such as file descriptors, threads, and mutexes, which can be affected by the use of the `fork` system call in `Ice::Service`. For example, any files opened in `main` are automatically closed by `Ice::Service` and therefore unusable in your service, unless the daemon is started with the `--noclose` option.

Windows Services

On Windows, `Ice::Service` recognizes the following command-line options:

- `--service` *NAME*
  Run as a Windows service named *NAME*, which must already be installed. This option is removed from the argument vector that is passed to the `start` member function.

Installing and configuring a Windows service is outside the scope of the `Ice::Service` class. Ice includes a utility for installing its services which you can use as a model for your own applications.

The `Ice::Service` class supports the Windows service control codes `SERVICE_CONTROL_INTERROGATE` and `SERVICE_CONTROL_STOP`. Upon receipt of `SERVICE_CONTROL_STOP`, `Ice::Service` invokes the `shutdown` member function.

`Ice::Service` Logging Considerations

A service that uses a custom logger has several ways of configuring it:

- as a process-wide logger,
- in the `InitializationData` argument that is passed to `main`,
- by overriding the `initializeCommunicator` member function.

On Windows, `Ice::Service` installs its own logger that uses the Windows `Application` event log if no custom logger is defined. The source name for the event log is the service's name unless a different value is specified using the property `Ice.EventLog.Source`.

On Unix, the default Ice logger (which logs to the standard error output) is used when no other logger is configured. For daemons, this is not appropriate because the output will be lost. To change this, you can either implement a custom logger or set the `Ice.UseSyslog` property, which selects a logger implementation that logs to the `syslog` facility. Alternatively, you can set the `Ice.LogFile` property to write log messages to a file.

Note that `Ice::Service` may encounter errors before the communicator is initialized. In this situation, `Ice::Service` uses its default logger unless a process-wide logger is configured. Therefore, even if a failing service is configured to use a different logger implementation, you may find useful diagnostic information in the `Application` event log (on Windows) or sent to standard error (on Unix).

See Also

- Hello World Application
- Properties and Configuration
- Communicator Initialization
- Logger Facility
- Portable Signal Handling in C++
- Windows Services

Copyright 2016, ZeroC, Inc.

Server-Side C++ Mapping for Interfaces

The server-side mapping for interfaces provides an up-call API for the Ice run time: by implementing virtual functions in a servant class, you provide the hook that gets the thread of control from the Ice server-side run time into your application code.

On this page:

- Skeleton Classes in C++
- Servant Classes in C++
  - Normal and idempotent Operations in C++

Skeleton Classes in C++

On the client side, interfaces map to proxy classes. On the server side, interfaces map to *skeleton* classes. A skeleton is a class that has a pure virtual member function for each operation on the corresponding interface. For example, consider our Slice definition for the `Node` interface:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };
    // ...
};
```

The Slice compiler generates the following definition for this interface:

C++

```
namespace Filesystem {

    class Node : public virtual Ice::Object {
    public:
        virtual std::string name(const Ice::Current& = Ice::Current()) = 0;
        // ...
    };
    // ...
}
```

For the moment, we will ignore a number of other member functions of this class. The important points to note are:

- As for the client side, Slice modules are mapped to C++ namespaces with the same name, so the skeleton class definition is nested in the namespace `Filesystem`.
- The name of the skeleton class is the same as the name of the Slice interface (`Node`).
- The skeleton class contains a pure virtual member function for each operation in the Slice interface.
- The skeleton class is an abstract base class because its member functions are pure virtual.
- The skeleton class inherits from `Ice::Object` (which forms the root of the Ice object hierarchy).

Servant Classes in C++

In order to provide an implementation for an Ice object, you must create a servant class that inherits from the corresponding skeleton class. For example, to create a servant for the `Node` interface, you could write:

Copyright 2016, ZeroC, Inc.

C++

```
#include <Filesystem.h> // Slice-generated header

class NodeI : public virtual Filesystem::Node {
public:
    NodeI(const std::string&);
    virtual std::string name(const Ice::Current&);
private:
    std::string _name;
};
```

By convention, servant classes have the name of their interface with an `I`-suffix, so the servant class for the `Node` interface is called `NodeI`. (This is a convention only: as far as the Ice run time is concerned, you can choose any name you prefer for your servant classes.)

Note that `NodeI` inherits from `Filesystem::Node`, that is, it derives from its skeleton class. It is a good idea to always use virtual inheritance when defining servant classes. Strictly speaking, virtual inheritance is necessary only for servants that implement interfaces that use multiple inheritance; however, the `virtual` keyword does no harm and, if you add multiple inheritance to an interface hierarchy half-way through development, you do not have to go back and add a `virtual` keyword to all your servant classes.

As far as Ice is concerned, the `NodeI` class must implement only a single member function: the pure virtual `name` function that it inherits from its skeleton. This makes the servant class a concrete class that can be instantiated. You can add other member functions and data members as you see fit to support your implementation. For example, in the preceding definition, we added a `_name` member and a constructor. Obviously, the constructor initializes the `_name` member and the `name` function returns its value:

C++

```
NodeI::NodeI(const std::string& name) : _name(name)
{
}

std::string
NodeI::name(const Ice::Current&) const
{
    return _name;
}
```

Normal and `idempotent` Operations in C++

The `name` member function of the `NodeI` skeleton is not a `const` member function. However, given that the operation does not modify the state of its object, it really should be a `const` member function. We can achieve this by adding the `["cpp:const"]` metadata directive. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Example {
                void normalOp();

    idempotent  void idempotentOp();

    ["cpp:const"]
    idempotent  void readonlyOp();
};
```

The skeleton class for this interface looks like this:

C++

```
class Example : public virtual Ice::Object {
public:
    virtual void normalOp(const Ice::Current& = Ice::Current()) = 0;
    virtual void idempotentOp(const Ice::Current& = Ice::Current()) = 0;
    virtual void readonlyOp(const Ice::Current& = Ice::Current()) const = 0;
    // ...
};
```

Note that `readonlyOp` is mapped as a `const` member function due to the `["cpp:const"]` metadata directive; normal and `idempotent` operations (without the metadata directive) are mapped as ordinary, non-`const` member functions.

See Also

- Slice for a Simple File System
- C++ Mapping for Interfaces
- Parameter Passing in C++
- Raising Exceptions in C++

Copyright 2016, ZeroC, Inc.

Parameter Passing in C++

For each parameter of a Slice operation, the C++ mapping generates a corresponding parameter for the virtual member function in the skeleton. In addition, every operation has a trailing parameter of type `Ice::Current`. For example, the `name` operation of the `Node` interface has no parameters, but the `name` member function of the `Node` skeleton class has a single parameter of type `Ice::Current`. We will ignore this parameter for now.

Parameter passing on the server side follows the rules for the client side:

- in-parameters are passed by value or `const` reference
- out-parameters are passed by reference
- return values are passed by value
- optional parameters are enclosed in `IceUtil::Optional` values

To illustrate the rules, consider the following interface that passes string parameters in all possible directions:

Slice

```
module M {
    interface Example {
        string op(string sin, out string sout);
    };
};
```

The generated skeleton class for this interface looks as follows:

C++

```
namespace M {
    class Example : virtual public ::Ice::Object {
    public:
        virtual std::string op(const std::string&, std::string&,
                               const Ice::Current& = Ice::Current()) = 0;
        // ...
    };
}
```

As you can see, there are no surprises here. For example, we could implement op as follows:

C++

```
std::string
ExampleI::op(const std::string& sin, std::string& sout, const Ice::Current&)
{
    cout << sin << endl;        // In parameters are initialized
    sout = "Hello World!";      // Assign out parameter
    return "Done";              // Return a string
}
```

This code is in no way different from what you would normally write if you were to pass strings to and from a function; the fact that remote

Copyright 2016, ZeroC, Inc.

procedure calls are involved does not impact on your code in any way. The same is true for parameters of other types, such as proxies, classes, or dictionaries: the parameter passing conventions follow normal C++ rules and do not require special-purpose API calls or memory management.

See Also

- Server-Side C++ Mapping for Interfaces
- C++ Mapping for Operations
- C++ Mapping for Optional Values
- Raising Exceptions in C++
- The Current Object

Copyright 2016, ZeroC, Inc.

Raising Exceptions in C++

To throw an exception from an operation implementation, you simply instantiate the exception, initialize it, and throw it. For example:

C++

```
void
Filesystem::FileI::write(const Filesystem::Lines& text, const Ice::Current&)
{
    // Try to write the file contents here...
    // Assume we are out of space...
    if (error) {
        Filesystem::GenericError e;
        e.reason = "file too large";
        throw e;
    }
};
```

No memory management issues arise in the presence of exceptions.

Note that the Slice compiler never generates exception specifications for operations, regardless of whether the corresponding Slice operation definition has an exception specification or not. This is deliberate: C++ exception specifications do not add any value and are therefore not used by the Ice C++ mapping [1].

If you throw an arbitrary C++ exception (such as an `int` or other unexpected type), the Ice run time catches the exception and then returns an `UnknownException` to the client. Similarly, if you throw an "impossible" user exception (a user exception that is not listed in the exception specification of the operation), the client receives an `UnknownUserException`.

If you throw a run-time exception, such as `MemoryLimitException`, the client receives an `UnknownLocalException`. For that reason, you should never throw system exceptions from operation implementations. If you do, all the client will see is an `UnknownLocalException`, which does not tell the client anything useful.

Three run-time exceptions are treated specially and not changed to `UnknownLocalException` when returned to the client: `ObjectNotExistException`, `OperationNotExistException`, and `FacetNotExistException`.

See Also

- Run-Time Exceptions
- C++ Mapping for Exceptions
- Server-Side C++ Mapping for Interfaces
- Parameter Passing in C++

References

1. Sutter, H. 2002. A Pragmatic Look at Exception Specifications. C/C++ Users Journal 20 (7): 59-64.

Copyright 2016, ZeroC, Inc.

Object Incarnation in C++

Having created a servant class such as the rudimentary `NodeI` class, you can instantiate the class to create a concrete servant that can receive invocations from a client. However, merely instantiating a servant class is insufficient to incarnate an object. Specifically, to provide an implementation of an Ice object, you must follow the following steps:

1. Instantiate a servant class.
2. Create an identity for the Ice object incarnated by the servant.
3. Inform the Ice run time of the existence of the servant.
4. Pass a proxy for the object to a client so the client can reach it.

On this page:

- Instantiating a C++ Servant
- Creating an Identity in C++
- Activating a C++ Servant
  - Servant Life Time and Reference Counts
- UUIDs as Identities in C++
- Creating Proxies in C++
  - Proxies and Servant Activation in C++
  - Direct Proxy Creation in C++

Instantiating a C++ Servant

Instantiating a servant means to allocate an instance on the heap:

C++

```
NodePtr servant = new NodeI("Fred");
```

This code creates a new `NodeI` instance on the heap and assigns its address to a smart pointer of type `NodePtr`. This works because `NodeI` is derived from `Node`, so a smart pointer of type `NodePtr` can also look after an instance of type `NodeI`. However, if we want to invoke a member function of the derived `NodeI` class at this point, we have a problem: we cannot access member functions of the derived `NodeI` class through a `NodePtr` smart pointer, only member functions of `Node` base class. (The C++ type rules prevent us from accessing a member of a derived class through a pointer to a base class.) To get around this, we can modify the code as follows:

C++

```
typedef IceUtil::Handle<NodeI> NodeIPtr;
NodeIPtr servant = new NodeI("Fred");
```

This code makes use of the smart pointer template by defining `NodeIPtr` as a smart pointer to `NodeI` instances. Whether you use a smart pointer of type `NodePtr` or `NodeIPtr` depends solely on whether you want to invoke a member function of the `NodeI` derived class; if you only want to invoke member functions that are defined in the `Node` skeleton base class, it is sufficient to use a `NodePtr` and you need not define the `NodeIPtr` type.

Whether you use `NodePtr` or `NodeIPtr`, the advantages of using a smart pointer class should be obvious from the smart pointer discussion: they make it impossible to accidentally leak memory.

Creating an Identity in C++

Each Ice object requires an identity. That identity must be unique for all servants using the same object adapter.

The Ice object model assumes that all objects (regardless of their adapter) have a globally unique identity.

An Ice object identity is a structure with the following Slice definition:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    struct Identity {
        string name;
        string category;
    };
    // ...
};
```

The full identity of an object is the combination of both the `name` and `category` fields of the `Identity` structure. For now, we will leave the `category` field as the empty string and simply use the `name` field. (The `category` field is most often used in conjunction with servant locators.)

To create an identity, we simply assign a key that identifies the servant to the `name` field of the `Identity` structure:

C++

```
Ice::Identity id;
id.name = "Fred"; // Not unique, but good enough for now
```

Activating a C++ Servant

Merely creating a servant instance does nothing: the Ice run time becomes aware of the existence of a servant only once you explicitly tell the object adapter about the servant. To activate a servant, you invoke the `add` operation on the object adapter. Assuming that we have access to the object adapter in the `_adapter` variable, we can write:

C++

```
_adapter->add(servant, id);
```

Note the two arguments to `add`: the smart pointer to the servant and the object identity. Calling `add` on the object adapter adds the servant pointer and the servant's identity to the adapter's servant map and links the proxy for an Ice object to the correct servant instance in the server's memory as follows:

1. The proxy for an Ice object, apart from addressing information, contains the identity of the Ice object. When a client invokes an operation, the object identity is sent with the request to the server.
2. The object adapter receives the request, retrieves the identity, and uses the identity as an index into the servant map.
3. If a servant with that identity is active, the object adapter retrieves the servant pointer from the servant map and dispatches the incoming request into the correct member function on the servant.

Assuming that the object adapter is in the active state, client requests are dispatched to the servant as soon as you call `add`.

Servant Life Time and Reference Counts

Putting the preceding points together, we can write a simple function that instantiates and activates one of our `NodeI` servants. For this example, we use a simple helper function called `activateServant` that creates and activates a servant with a given identity:

Copyright 2016, ZeroC, Inc.

C++

```
void
activateServant(const string& name)
{
    NodePtr servant = new NodeI(name);          // Refcount == 1
    Ice::Identity id;
    id.name = name;
    _adapter->add(servant, id);                 // Refcount == 2
}                                               // Refcount == 1
```

Note that we create the servant on the heap and that, once `activateServant` returns, we lose the last remaining handle to the servant (because the `servant` variable goes out of scope). The question is, what happens to the heap-allocated servant instance? The answer lies in the smart pointer semantics:

- When the new servant is instantiated, its reference count is initialized to 0.
- Assigning the servant's address to the `servant` smart pointer increments the servant's reference count to 1.
- Calling `add` passes the `servant` smart pointer to the object adapter which keeps a copy of the handle internally. This increments the reference count of the servant to 2.
- When `activateServant` returns, the destructor of the `servant` variable decrements the reference count of the servant to 1.

The net effect is that the servant is retained on the heap with a reference count of 1 for as long as the servant is in the servant map of its object adapter. (If we deactivate the servant, that is, remove it from the servant map, the reference count drops to zero and the memory occupied by the servant is reclaimed; we discuss these life cycle issues in Object Life Cycle.)

UUIDs as Identities in C++

The Ice object model assumes that object identities are globally unique. One way of ensuring that uniqueness is to use UUIDs (Universally Unique Identifiers) as identities. The `IceUtil` namespace contains a helper function to create such identities:

C++

```
#include <IceUtil/UUID.h>
#include <iostream>

using namespace std;

int
main()
{
    cout << IceUtil::generateUUID() << endl;
}
```

When executed, this program prints a unique string such as `5029a22c-e333-4f87-86b1-cd5e0fcce509`. Each call to `generateUUID` creates a string that differs from all previous ones.

You can use a UUID such as this to create object identities. For convenience, the object adapter has an operation `addWithUUID` that generates a UUID and adds a servant to the servant map in a single step. Using this operation, we can rewrite the code shown earlier like this:

Copyright 2016, ZeroC, Inc.

C++

```
void
activateServant(const string& name)
{
    NodePtr servant = new NodeI(name);
    _adapter->addWithUUID(servant);
}
```

Creating Proxies in C++

Once we have activated a servant for an Ice object, the server can process incoming client requests for that object. However, clients can only access the object once they hold a proxy for the object. If a client knows the server's address details and the object identity, it can create a proxy from a string, as we saw in our first example in Hello World Application. However, creation of proxies by the client in this manner is usually only done to allow the client access to initial objects for bootstrapping. Once the client has an initial proxy, it typically obtains further proxies by invoking operations.

The object adapter contains all the details that make up the information in a proxy: the addressing and protocol information, and the object identity. The Ice run time offers a number of ways to create proxies. Once created, you can pass a proxy to the client as the return value or as an out-parameter of an operation invocation.

Proxies and Servant Activation in C++

The `add` and `addWithUUID` servant activation operations on the object adapter return a proxy for the corresponding Ice object. This means we can write:

C++

```
typedef IceUtil::Handle<NodeI> NodeIPtr;
NodeIPtr servant = new NodeI(name);
NodePrx proxy = NodePrx::uncheckedCast(_adapter->addWithUUID(servant));

// Pass proxy to client...
```

Here, `addWithUUID` both activates the servant and returns a proxy for the Ice object incarnated by that servant in a single step.

Note that we need to use an `uncheckedCast` here because `addWithUUID` returns a proxy of type `Ice::ObjectPrx`.

Direct Proxy Creation in C++

The object adapter offers an operation to create a proxy for a given identity:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local interface ObjectAdapter {
        Object* createProxy(Identity id);
        // ...
    };
};
```

Note that `createProxy` creates a proxy for a given identity whether a servant is activated with that identity or not. In other words, proxies have a life cycle that is quite independent from the life cycle of servants:

C++

```
Ice::Identity id;
id.name = IceUtil::generateUUID();
ObjectPrx o = _adapter->createProxy(id);
```

This creates a proxy for an Ice object with the identity returned by `generateUUID`. Obviously, no servant yet exists for that object so, if we return the proxy to a client and the client invokes an operation on the proxy, the client will receive an `ObjectNotExistException`. (We examine these life cycle issues in more detail in Object Life Cycle.)

See Also

- Hello World Application
- Object Adapter States
- Servant Locators
- Object Life Cycle
- The C++ generateUUID Function

Copyright 2016, ZeroC, Inc.

Asynchronous Method Dispatch (AMD) in C++

The number of simultaneous synchronous requests a server is capable of supporting is determined by the number of threads in the server's thread pool. If all of the threads are busy dispatching long-running operations, then no threads are available to process new requests and therefore clients may experience an unacceptable lack of responsiveness.

Asynchronous Method Dispatch (AMD), the server-side equivalent of AMI, addresses this scalability issue. Using AMD, a server can receive a request but then suspend its processing in order to release the dispatch thread as soon as possible. When processing resumes and the results are available, the server sends a response explicitly using a callback object provided by the Ice run time.

AMD is transparent to the client, that is, there is no way for a client to distinguish a request that, in the server, is processed synchronously from a request that is processed asynchronously.

In practical terms, an AMD operation typically queues the request data (i.e., the callback object and operation arguments) for later processing by an application thread (or thread pool). In this way, the server minimizes the use of dispatch threads and becomes capable of efficiently supporting thousands of simultaneous clients.

An alternate use case for AMD is an operation that requires further processing after completing the client's request. In order to minimize the client's delay, the operation returns the results while still in the dispatch thread, and then continues using the dispatch thread for additional work.

On this page:

- Enabling AMD with Metadata in C++
- AMD Mapping in C++
- AMD Exceptions in C++
- AMD Example in C++

Enabling AMD with Metadata in C++

To enable asynchronous dispatch, you must add an `["amd"]` metadata directive to your Slice definitions. The directive applies at the interface and the operation level. If you specify `["amd"]` at the interface level, all operations in that interface use asynchronous dispatch; if you specify `["amd"]` for an individual operation, only that operation uses asynchronous dispatch. In either case, the metadata directive *replaces* synchronous dispatch, that is, a particular operation implementation must use synchronous or asynchronous dispatch and cannot use both.

Consider the following Slice definitions:

Slice

```
["amd"] interface I {
    bool isValid();
    float computeRate();
};

interface J {
    ["amd"] void startProcess();
    int endProcess();
};
```

In this example, both operations of interface `I` use asynchronous dispatch, whereas, for interface `J`, `startProcess` uses asynchronous dispatch and `endProcess` uses synchronous dispatch.

Specifying metadata at the operation level (rather than at the interface or class level) minimizes the amount of generated code and, more importantly, minimizes complexity: although the asynchronous model is more flexible, it is also more complicated to use. It is therefore in your best interest to limit the use of the asynchronous model to those operations that need it, while using the simpler synchronous model for the rest.

AMD Mapping in C++

The C++ mapping emits the following code for each AMD operation:

Copyright 2016, ZeroC, Inc.

1. A callback class used by the implementation to notify the Ice run time about the completion of an operation. The name of this class is formed using the pattern `AMD_`*`class_op`*. For example, an operation named `foo` defined in interface `I` results in a class named `AMD_I_foo`. The class is generated in the same scope as the interface or class containing the operation. Several methods are provided:
   - `void ice_response(<params>);`
     `ice_responsevoidice_responseout`
   - `void ice_exception(const std::exception &);`
     `ice_exception`
   - `void ice_exception();`
     `ice_exceptionUnknownException`

   Neither `ice_response` nor `ice_exception` throw any exceptions to the caller.
2. The dispatch method, whose name has the suffix `_async`. This method has a `void` return type. The first parameter is a smart pointer to an instance of the callback class described above. The remaining parameters comprise the in-parameters of the operation, in the order of declaration.

For example, suppose we have defined the following operation:

Slice

```
interface I {
    ["amd"] int foo(short s, out long l);
};
```

The callback class generated for operation `foo` is shown below:

C++

```
class AMD_I_foo : public ... {
public:
    void ice_response(Ice::Int, Ice::Long);
    void ice_exception(const std::exception&);
    void ice_exception();
};
```

The dispatch method for asynchronous invocation of operation `foo` is generated as follows:

C++

```
void foo_async(const AMD_I_fooPtr&, Ice::Short);
```

AMD Exceptions in C++

There are two processing contexts in which the logical implementation of an AMD operation may need to report an exception: the dispatch thread (the thread that receives the invocation), and the response thread (the thread that sends the response).

These are not necessarily two different threads: it is legal to send the response from the dispatch thread.

Although we recommend that the callback object be used to report all exceptions to the client, it is legal for the implementation to raise an exception instead, but only from the dispatch thread.

Copyright 2016, ZeroC, Inc.

As you would expect, an exception raised from a response thread cannot be caught by the Ice run time; the application's run-time environment determines how such an exception is handled. Therefore, a response thread must ensure that it traps all exceptions and sends the appropriate response using the callback object. Otherwise, if a response thread is terminated by an uncaught exception, the request may never be completed and the client might wait indefinitely for a response.

Whether raised in a dispatch thread or reported via the callback object, user exceptions are validated and local exceptions may undergo translation.

AMD Example in C++

To demonstrate the use of AMD in Ice, let us define the Slice interface for a simple computational engine:

Slice

```
module Demo {
    sequence<float> Row;
    sequence<Row> Grid;

    exception RangeError {};

    interface Model {
        ["amd"] Grid interpolate(Grid data, float factor)
            throws RangeError;
    };
};
```

Given a two-dimensional grid of floating point values and a factor, the `interpolate` operation returns a new grid of the same size with the values interpolated in some interesting (but unspecified) way.

Our servant class derives from `Demo::Model` and supplies a definition for the `interpolate_async` method:

C++

```
class ModelI : virtual public Demo::Model, virtual public IceUtil::Mutex {
public:
    virtual void interpolate_async(
        const Demo::AMD_Model_interpolatePtr&,
        const Demo::Grid&,
        Ice::Float,
        const Ice::Current&);

private:
  std::list<JobPtr> _jobs;
};
```

The implementation of `interpolate_async` uses synchronization to safely record the callback object and arguments in a `Job` that is added to a queue:

Copyright 2016, ZeroC, Inc.

C++

```
void ModelI::interpolate_async(
    const Demo::AMD_Model_interpolatePtr& cb,
    const Demo::Grid& data,
    Ice::Float factor,
    const Ice::Current& current)
{
    IceUtil::Mutex::Lock sync(*this);
    JobPtr job = new Job(cb, data, factor);
    _jobs.push_back(job);
}
```

After queuing the information, the operation returns control to the Ice run time, making the dispatch thread available to process another request. An application thread removes the next `Job` from the queue and invokes `execute` to perform the interpolation. `Job` is defined as follows:

C++

```
class Job : public IceUtil::Shared {
public:
    Job(const Demo::AMD_Model_interpolatePtr&, const Demo::Grid&, Ice::Float);
    void execute();

private:
    bool interpolateGrid();

    Demo::AMD_Model_interpolatePtr _cb;
    Demo::Grid _grid;
    Ice::Float _factor;
};
typedef IceUtil::Handle<Job> JobPtr;
```

The implementation of `execute` uses `interpolateGrid` (not shown) to perform the computational work:

Copyright 2016, ZeroC, Inc.

C++

```
Job::Job(const Demo::AMD_Model_interpolatePtr& cb, const Demo::Grid& grid,
Ice::Float factor) :
    _cb(cb), _grid(grid), _factor(factor)
{
}

void Job::execute()
{
    if (!interpolateGrid()) {
        _cb->ice_exception(Demo::RangeError());
        return;
    }
    _cb->ice_response(_grid);
}
```

If `interpolateGrid` returns `false`, then `ice_exception` is invoked to indicate that a range error has occurred. The `return` statement following the call to `ice_exception` is necessary because `ice_exception` does not throw an exception; it only marshals the exception argument and sends it to the client.

If interpolation was successful, `ice_response` is called to send the modified grid back to the client.

See Also

- Asynchronous Method Invocation (AMI) in C++
- The Ice Threading Model
- User Exceptions
- Run-Time Exceptions

Copyright 2016, ZeroC, Inc.

Example of a File System Server in C++

This page presents the source code for a C++ server that implements our file system and communicates with the client we wrote earlier. The code is fully functional, apart from the required interlocking for threads.

The server is remarkably free of code that relates to distribution: most of the server code is simply application logic that would be present just the same for a non-distributed version. Again, this is one of the major advantages of Ice: distribution concerns are kept away from application code so that you can concentrate on developing application logic instead of networking infrastructure.

The server code shown here is not quite correct as it stands: if two clients access the same file in parallel, each via a different thread, one thread may read the `_lines` data member while another thread updates it. Obviously, if that happens, we may write or return garbage or, worse, crash the server. However, it is trivial to make the `read` and `write` operations thread-safe: a single data member and two lines of source code are sufficient to achieve this. We discuss how to write thread-safe servant implementations in Threads and Concurrency with C++.

On this page:

- Implementing a File System Server in C++
- Server main Program in C++
- Servant Class Definitions in C++
- The Servant Implementation in C++
  - Implementing FileI
  - Implementing DirectoryI
  - Implementing NodeI

Implementing a File System Server in C++

We have now seen enough of the server-side C++ mapping to implement a server for our file system. (You may find it useful to review these Slice definitions before studying the source code.)

Our server is composed of two source files:

- `Server.cpp`
  This file contains the server main program.
- `FilesystemI.cpp`
  This file contains the implementation for the file system servants.

Server `main` Program in C++

Our server main program, in the file `Server.cpp`, uses the `Ice::Application` class. The `run` method installs a signal handler, creates an object adapter, instantiates a few servants for the directories and files in the file system, and then activates the adapter. This leads to a `main` program as follows:

C++

```
#include <Ice/Ice.h>
#include <FilesystemI.h>

using namespace std;
using namespace Filesystem;

class FilesystemApp : virtual public Ice::Application {
public:
    virtual int run(int, char*[]) {
        // Terminate cleanly on receipt of a signal
        //
        shutdownOnInterrupt();

        // Create an object adapter.
        //
```


Copyright 2016, ZeroC, Inc.

```
        Ice::ObjectAdapterPtr adapter =
communicator()->createObjectAdapterWithEndpoints(
            "SimpleFilesystem", "default -p 10000");

        // Create the root directory (with name "/" and no parent)
        //
        DirectoryIPtr root = new DirectoryI(communicator(), "/", 0);
        root->activate(adapter);

        // Create a file called "README" in the root directory
        //
        FileIPtr file = new FileI(communicator(), "README", root);
        Lines text;
        text.push_back("This file system contains a collection of poetry.");
        file->write(text);
        file->activate(adapter);

        // Create a directory called "Coleridge"
        // in the root directory
        //
        DirectoryIPtr coleridge =
new DirectoryI(communicator(), "Coleridge", root);
        coleridge->activate(adapter);

        // Create a file called "Kubla_Khan"
        // in the Coleridge directory
        //
        file = new FileI(communicator(), "Kubla_Khan", coleridge);
        text.erase(text.begin(), text.end());
        text.push_back("In Xanadu did Kubla Khan");
        text.push_back("A stately pleasure-dome decree:");
        text.push_back("Where Alph, the sacred river, ran");
        text.push_back("Through caverns measureless to man");
        text.push_back("Down to a sunless sea.");
        file->write(text);
        file->activate(adapter);

        // All objects are created, allow client requests now
        //
        adapter->activate();

        // Wait until we are done
        //
        communicator()->waitForShutdown();
        if (interrupted()) {
            cerr << appName() << ": received signal, shutting down" << endl;
        }

        return 0;
    };
};

int
main(int argc, char* argv[])
{
    FilesystemApp app;
```


Copyright 2016, ZeroC, Inc.

```
    return app.main(argc, argv);
}
```

There is quite a bit of code here, so let us examine each section in detail:

C++

```
#include <FilesystemI.h>
#include <Ice/Application.h>

using namespace std;
using namespace Filesystem;
```

The code includes the header file `FilesystemI.h`. That file includes `Ice/Ice.h` as well as the header file that is generated by the Slice compiler, `Filesystem.h`. Because we are using `Ice::Application`, we need to include `Ice/Application.h` as well.

Two `using` declarations, for the namespaces `std` and `Filesystem`, permit us to be a little less verbose in the source code.

The next part of the source code is the definition of `FilesystemApp`, which derives from `Ice::Application` and contains the main application logic in its `run` method:

C++

```
class FilesystemApp : virtual public Ice::Application {
public:
    virtual int run(int, char*[]) {
        // Terminate cleanly on receipt of a signal
        //
        shutdownOnInterrupt();

        // Create an object adapter.
        //
        Ice::ObjectAdapterPtr adapter =
communicator()->createObjectAdapterWithEndpoints(
            "SimpleFilesystem", "default -p 10000");

        // Create the root directory (with name "/" and no parent)
        //
        DirectoryIPtr root = new DirectoryI(communicator(), "/", 0);
        root->activate(adapter);

        // Create a file called "README" in the root directory
        //
        FileIPtr file = new FileI(communicator(), "README", root);
        Lines text;
        text.push_back("This file system contains a collection of poetry.");
        file->write(text);
        file->activate(adapter);

        // Create a directory called "Coleridge"
        // in the root directory
        //
        DirectoryIPtr coleridge =
```


Copyright 2016, ZeroC, Inc.

```
new DirectoryI(communicator(), "Coleridge", root);
        coleridge->activate(adapter);

        // Create a file called "Kubla_Khan"
        // in the Coleridge directory
        //
        file = new FileI(communicator(), "Kubla_Khan", coleridge);
        text.erase(text.begin(), text.end());
        text.push_back("In Xanadu did Kubla Khan");
        text.push_back("A stately pleasure-dome decree:");
        text.push_back("Where Alph, the sacred river, ran");
        text.push_back("Through caverns measureless to man");
        text.push_back("Down to a sunless sea.");
        file->write(text);
        file->activate(adapter);

        // All objects are created, allow client requests now
        //
        adapter->activate();

        // Wait until we are done
        //
        communicator()->waitForShutdown();
        if (interrupted()) {
            cerr << appName() << ": received signal, shutting down" << endl;
        }

        return 0;
```


Copyright 2016, ZeroC, Inc.

```
    };
};
```

Much of this code is boiler plate that we saw previously: we create an object adapter, and, towards the end, activate the object adapter and call `waitForShutdown`.

The interesting part of the code follows the adapter creation: here, the server instantiates a few nodes for our file system to create the structure shown below:

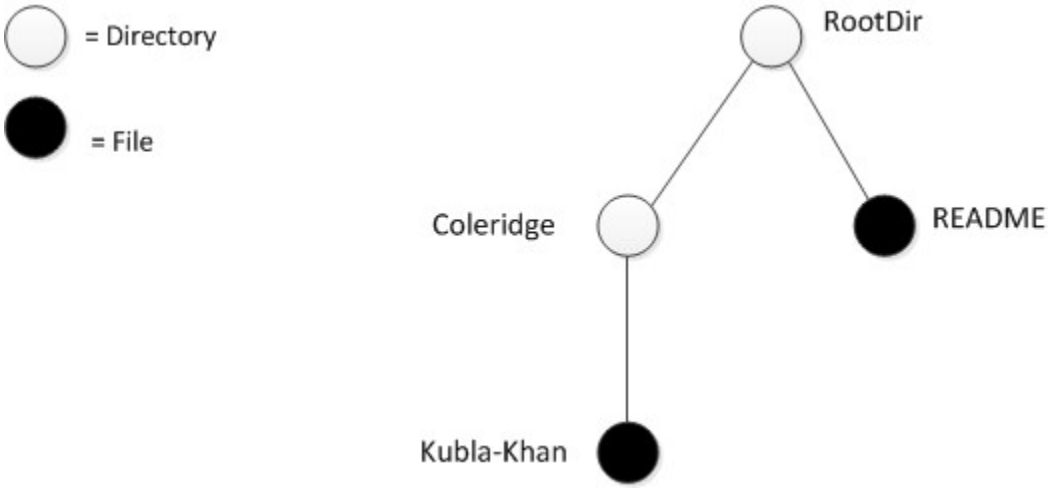

A small file system.

As we will see shortly, the servants for our directories and files are of type `DirectoryI` and `FileI`, respectively. The constructor for either type of servant accepts three parameters: the communicator, the name of the directory or file to be created, and a handle to the servant for the parent directory. (For the root directory, which has no parent, we pass a null parent handle.) Thus, the statement

C++

```
DirectoryIPtr root = new DirectoryI(communicator(), "/", 0);
```

creates the root directory, with the name `"/"` and no parent directory. Note that we use the smart pointer class to hold the return value from `new`; that way, we avoid any memory management issues. The types `DirectoryIPtr` and `FileIPtr` are defined as follows in a header file `FilesystemI.h`:

C++

```
typedef IceUtil::Handle<DirectoryI> DirectoryIPtr;
typedef IceUtil::Handle<FileI> FileIPtr;
```

Here is the code that establishes the structure in the illustration above.

Copyright 2016, ZeroC, Inc.

C++

```
        // Create the root directory (with name "/" and no parent)
        //
        DirectoryIPtr root = new DirectoryI(communicator(), "/", 0);
        root->activate(adapter);

        // Create a file called "README" in the root directory
        //
        FileIPtr file = new FileI(communicator(), "README", root);
        Lines text;
        text.push_back("This file system contains a collection of poetry.");
        file->write(text);
        file->activate(adapter);

        // Create a directory called "Coleridge"
        // in the root directory
        //
        DirectoryIPtr coleridge =
new DirectoryI(communicator(), "Coleridge", root);
        coleridge->activate(adapter);

        // Create a file called "Kubla_Khan"
        // in the Coleridge directory
        //
        file = new FileI(communicator(), "Kubla_Khan", coleridge);
        text.erase(text.begin(), text.end());
        text.push_back("In Xanadu did Kubla Khan");
        text.push_back("A stately pleasure-dome decree:");
        text.push_back("Where Alph, the sacred river, ran");
        text.push_back("Through caverns measureless to man");
        text.push_back("Down to a sunless sea.");
        file->write(text);
        file->activate(adapter);
```

We first create the root directory and a file `README` within the root directory. (Note that we pass the handle to the root directory as the parent pointer when we create the new node of type `FileI`.)

After creating each servant, the code calls `activate` on the servant. (We will see the definition of this member function shortly.) The `activate` member function adds the servant to the ASM.

The next step is to fill the file with text:

C++

```
        FileIPtr file = new FileI(communicator(), "README", root);
        Lines text;
        text.push_back("This file system contains a collection of poetry.");
        file->write(text);
        file->activate(adapter);
```

Recall that Slice sequences map to STL vectors. The Slice type `Lines` is a sequence of strings, so the C++ type `Lines` is a vector of strings; we add a line of text to our `README` file by calling `push_back` on that vector.

Copyright 2016, ZeroC, Inc.

Finally, we call the Slice `write` operation on our `FileI` servant by simply writing:

C++

```
        file->write(text);
```

This statement is interesting: the server code invokes an operation on one of its own servants. Because the call happens via a smart class pointer (of type `FilePtr`) and not via a proxy (of type `FilePrx`), the Ice run time does not know that this call is even taking place — such a direct call into a servant is not mediated by the Ice run time in any way and is dispatched as an ordinary C++ function call.

In similar fashion, the remainder of the code creates a subdirectory called `Coleridge` and, within that directory, a file called `Kubla_Khan` to complete the structure in the above illustration.

Servant Class Definitions in C++

We must provide servants for the concrete interfaces in our Slice specification, that is, we must provide servants for the `File` and `Directory` interfaces in the C++ classes `FileI` and `DirectoryI`. This means that our servant classes might look as follows:

C++

```
namespace Filesystem {
  class FileI : virtual public File {
    // ...
  };

  class DirectoryI : virtual public Directory {
     // ...
  };
}
```

This leads to the C++ class structure as shown:

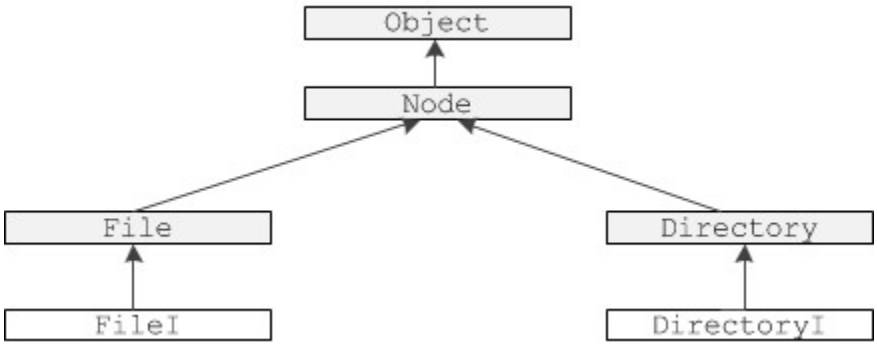

File system servants using interface inheritance.

The shaded classes in the illustration above are skeleton classes and the unshaded classes are our servant implementations. If we implement our servants like this, `FileI` must implement the pure virtual operations it inherits from the `File` skeleton (`read` and `write`), as well as the operation it inherits from the `Node` skeleton (`name`). Similarly, `DirectoryI` must implement the pure virtual function it inherits from the `Directory` skeleton (`list`), as well as the operation it inherits from the `Node` skeleton (`name`). Implementing the servants in this way uses interface inheritance from `Node` because no implementation code is inherited from that class.

Alternatively, we can implement our servants using the following definitions:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Filesystem {
  class NodeI : virtual public Node {
    // ...
  };

  class FileI : virtual public File, virtual public NodeI {
    // ...
  };

  class DirectoryI : virtual public Directory, virtual public NodeI {
    // ...
  };
}
```

This leads to the C++ class structure shown:

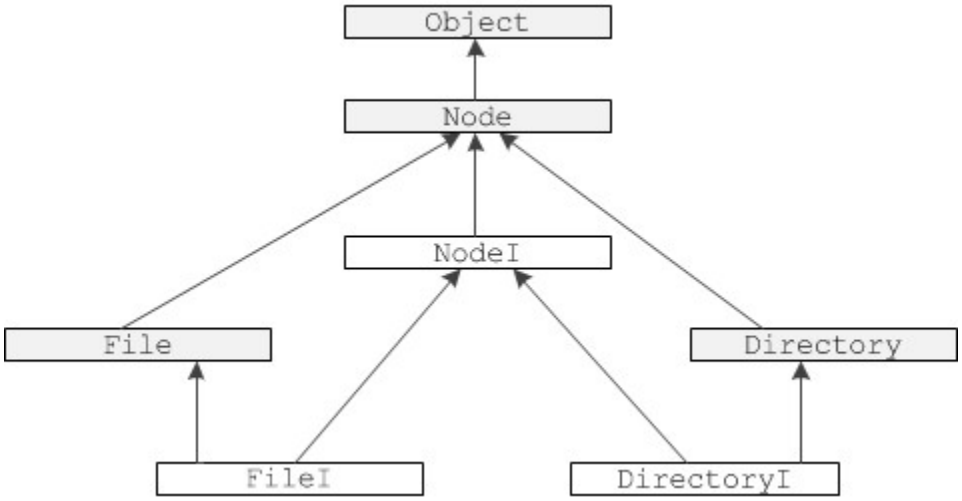

File system servants using implementation inheritance.

In this implementation, `NodeI` is a concrete base class that implements the `name` operation it inherits from the `Node` skeleton. `FileI` and `DirectoryI` use multiple inheritance from `NodeI` and their respective skeletons, that is, `FileI` and `DirectoryI` use implementation inheritance from their `NodeI` base class.

Either implementation approach is equally valid. Which one to choose simply depends on whether we want to re-use common code provided by `NodeI`. For the implementation that follows, we have chosen the second approach, using implementation inheritance.

Given the structure in the above illustration and the operations we have defined in the Slice definition for our file system, we can add these operations to the class definition for our servants:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Filesystem {
  class NodeI : virtual public Node {
  public:
    virtual std::string name(const Ice::Current&);
  };

  class FileI : virtual public File, virtual public NodeI {
  public:
    virtual Lines read(const Ice::Current&);
    virtual void write(const Lines&, const Ice::Current&);
  };

  class DirectoryI : virtual public Directory, virtual public NodeI {
  public:
    virtual NodeSeq list(const Ice::Current&);
  };
}
```

This simply adds signatures for the operation implementations to each class. Note that the signatures must exactly match the operation signatures in the generated skeleton classes — if they do not match exactly, you end up overloading the pure virtual function in the base class instead of overriding it, meaning that the servant class cannot be instantiated because it will still be abstract. To avoid signature mismatches, you can copy the signatures from the generated header file (`Filesystem.h`), or you can use the `--impl` option with `slice2cpp` to generate header and implementation files that you can add your application code to.

Now that we have the basic structure in place, we need to think about other methods and data members we need to support our servant implementation. Typically, each servant class hides the copy constructor and assignment operator, and has a constructor to provide initial state for its data members. Given that all nodes in our file system have both a name and a parent directory, this suggests that the `NodeI` class should implement the functionality relating to tracking the name of each node, as well as the parent-child relationships:

C++

```
namespace Filesystem {
  class DirectoryI;
  typedef IceUtil::Handle<DirectoryI> DirectoryIPtr;

  class NodeI : virtual public Node {
  public:
    virtual std::string name(const Ice::Current&);
    NodeI(const Ice::CommunicatorPtr&, const std::string&, const DirectoryIPtr&);
    void activate(const Ice::ObjectAdapterPtr&);
  private:
    std::string _name;
    Ice::Identity _id;
    DirectoryIPtr _parent;
    NodeI(const NodeI&);               // Copy forbidden
    void operator=(const NodeI&);      // Assignment forbidden
  };
}
```

The `NodeI` class has a private data member to store its name (of type `std::string`) and its parent directory (of type `DirectoryIPtr`). The constructor accepts parameters that set the value of these data members. For the root directory, by convention, we pass a null handle to

Copyright 2016, ZeroC, Inc.

the constructor to indicate that the root directory has no parent. The constructor also requires the communicator to be passed to it. This is necessary because the constructor creates the identity for the servant, which requires access to the communicator. The `activate` member function adds the servant to the ASM (which requires access to the object adapter) and connects the child to its parent.

The `FileI` servant class must store the contents of its file, so it requires a data member for this. We can conveniently use the generated `Lines` type (which is a `std::vector<std::string>`) to hold the file contents, one string for each line. Because `FileI` inherits from `NodeI`, it also requires a constructor that accepts the communicator, file name, and parent directory, leading to the following class definition:

C++

```
namespace Filesystem {
  class FileI : virtual public File, virtual public NodeI {
  public:
    virtual Lines read(const Ice::Current&);
    virtual void write(const Lines&, const Ice::Current&);
    FileI(const Ice::CommunicatorPtr&, const std::string&, const DirectoryIPtr&);
  private:
    Lines _lines;
  };
}
```

For directories, each directory must store its list of child notes. We can conveniently use the generated `NodeSeq` type (which is a `vector<NodePrx>`) to do this. Because `DirectoryI` inherits from `NodeI`, we need to add a constructor to initialize the directory name and its parent directory. As we will see shortly, we also need a private helper function, `addChild`, to make it easier to connect a newly created directory to its parent. This leads to the following class definition:

C++

```
namespace Filesystem {
  class DirectoryI : virtual public Directory, virtual public NodeI {
  public:
    virtual NodeSeq list(const Ice::Current&) const;
    DirectoryI(const Ice::CommunicatorPtr&, const std::string&,
const DirectoryIPtr&);
    void addChild(NodePrx child);
  private:
    NodeSeq _contents;
  };
}
```

Servant Header File Example

Putting all this together, we end up with a servant header file, `FilesystemI.h`, as follows:

Copyright 2016, ZeroC, Inc.

C++

```
#include <Ice/Ice.h>
#include <Filesystem.h>

namespace Filesystem {
  class DirectoryI;
  typedef IceUtil::Handle<DirectoryI> DirectoryIPtr;

  class NodeI : virtual public Node {
  public:
    virtual std::string name(const Ice::Current&);
    NodeI(const Ice::CommunicatorPtr&, const std::string&, const DirectoryIPtr&);
    void activate(const Ice::ObjectAdapterPtr&);
  private:
    std::string _name;
    Ice::Identity _id;
    DirectoryIPtr _parent;
    NodeI(const NodeI&);            // Copy forbidden
    void operator=(const NodeI&);   // Assignment forbidden
  };

  typedef IceUtil::Handle<NodeI> NodeIPtr;

  class FileI : virtual public File, virtual public NodeI {
  public:
    virtual Lines read(const Ice::Current&);
    virtual void write(const Lines&, const Ice::Current& = Ice::Current());
    FileI(const Ice::CommunicatorPtr&, const std::string&, const DirectoryIPtr&);
  private:
    Lines _lines;
  };

  typedef IceUtil::Handle<FileI> FileIPtr;

  class DirectoryI : virtual public Directory, virtual public NodeI {
  public:
    virtual NodeSeq list(const Ice::Current&);
    DirectoryI(const Ice::CommunicatorPtr&, const std::string&,
const DirectoryIPtr&);
    void addChild(const Filesystem::NodePrx&);
  private:
    Filesystem::NodeSeq _contents;
  };
}
```

The Servant Implementation in C++

The implementation of our servants is mostly trivial, following from the class definitions in our `FilesystemI.h` header file.

Implementing `FileI`

The implementation of the `read` and `write` operations for files is trivial: we simply store the passed file contents in the `_lines` data

Copyright 2016, ZeroC, Inc.

member. The constructor is equally trivial, simply passing its arguments through to the `NodeI` base class constructor:

C++

```
Filesystem::Lines
Filesystem::FileI::read(const Ice::Current&)
{
    return _lines;
}

void
Filesystem::FileI::write(const Filesystem::Lines& text, const Ice::Current&)
{
    _lines = text;
}

Filesystem::FileI::FileI(const Ice::CommunicatorPtr& communicator,
                         const string& name,
                         const DirectoryIPtr& parent)
    : NodeI(communicator, name, parent)
{
}
```

Implementing `DirectoryI`

The implementation of `DirectoryI` is equally trivial: the `list` operation simply returns the `_contents` data member and the constructor passes its arguments through to the `NodeI` base class constructor:

C++

```
Filesystem::NodeSeq
Filesystem::DirectoryI::list(const Ice::Current&)
{
    return _contents;
}

Filesystem::DirectoryI::DirectoryI(const Ice::CommunicatorPtr& communicator,
                                   const string& name,
                                   const DirectoryIPtr& parent)
    : NodeI(name, parent)
{
}

void
Filesystem::DirectoryI::addChild(const NodePrx child)
{
    _contents.push_back(child);
}
```

The only noteworthy thing is the implementation of `addChild`: when a new directory or file is created, the constructor of the `NodeI` base class calls `addChild` on its own parent, passing it the proxy to the newly-created child. The implementation of `addChild` appends the passed reference to the contents list of the directory it is invoked on (which is the parent directory).

Copyright 2016, ZeroC, Inc.

Implementing `NodeI`

The name operation of our `NodeI` class is again trivial: it simply returns the `_name` data member:

C++

```
std::string
Filesystem::NodeI::name(const Ice::Current&)
{
    return _name;
}
```

The `NodeI` constructor creates an identity for the servant:

C++

```
Filesystem::NodeI::NodeI(const Ice::CommunicatorPtr& communicator,
                         const string& name,
                         const DirectoryIPtr& parent)
    : _name(name), _parent(parent)
{
    _id.name = parent ? IceUtil::generateUUID() : "RootDir";
}
```

For the root directory, we use the fixed identity `"RootDir"`. This allows the client to create a proxy for the root directory. For directories other than the root directory, we use a UUID as the identity.

Finally, `NodeI` provides the `activate` member function that adds the servant to the ASM and connects the child node to its parent directory:

C++

```
void
Filesystem::NodeI::activate(const Ice::ObjectAdapterPtr& a)
{
    NodePrx thisNode = NodePrx::uncheckedCast(a->add(this, _id));
    if(_parent)
    {
        _parent->addChild(thisNode);
    }
}
```

This completes our servant implementation. The complete source code is shown here once more:

C++

```
#include <IceUtil/IceUtil.h>
#include <FilesystemI.h>

using namespace std;
```


Copyright 2016, ZeroC, Inc.

```
// Slice Node::name() operation

std::string
Filesystem::NodeI::name(const Ice::Current&)
{
    return _name;
}

// NodeI constructor

Filesystem::NodeI::NodeI(const Ice::CommunicatorPtr& communicator,
                         const string& name,
                         const DirectoryIPtr& parent)
    : _name(name), _parent(parent)
{
    // Create an identity. The root directory has the fixed identity "RootDir"
    //
    _id.name = parent ? IceUtil::generateUUID() : "RootDir";
}

// NodeI activate() member function

void
Filesystem::NodeI::activate(const Ice::ObjectAdapterPtr& a)
{
    NodePrx thisNode = NodePrx::uncheckedCast(a->add(this, _id));
    if(_parent)
    {
        _parent->addChild(thisNode);
    }
}

// Slice File::read() operation

Filesystem::Lines
Filesystem::FileI::read(const Ice::Current&)
{
    return _lines;
}

// Slice File::write() operation

void
Filesystem::FileI::write(const Filesystem::Lines& text, const Ice::Current&)
{
    _lines = text;
}

// FileI constructor

Filesystem::FileI::FileI(const Ice::CommunicatorPtr& communicator,
                         const string& name,
                         const DirectoryIPtr& parent)
    : NodeI(communicator, name, parent)
{
}
```


Copyright 2016, ZeroC, Inc.

```
// Slice Directory::list() operation

Filesystem::NodeSeq
Filesystem::DirectoryI::list(const Ice::Current& c)
{
    return _contents;
}

// DirectoryI constructor

Filesystem::DirectoryI::DirectoryI(const Ice::CommunicatorPtr& communicator,
                                   const string& name,
                                   const DirectoryIPtr& parent)
    : NodeI(communicator, name, parent)
{
}

// addChild is called by the child in order to add
// itself to the _contents member of the parent

void
Filesystem::DirectoryI::addChild(const NodePrx& child)
{
```


Copyright 2016, ZeroC, Inc.

```
    _contents.push_back(child);
}
```

See Also

- Slice for a Simple File System
- Example of a File System Client in C++
- The `Ice::Application` Class
- C++ Mapping for Sequences
- slice2cpp Command-Line Options
- UUIDs as Identities in C++
- Threads and Concurrency with C++

Copyright 2016, ZeroC, Inc.

Customizing the C++ Mapping

Ice for C++ allows you to map Slice strings, sequences and dictionaries to your own C++ types, through the `cpp:type` and `cpp:view-type` metadata directives.

When you map a Slice type to a C++ type from the standard library type, Ice is able to marshal and unmarshal this type without your help. For example, if you map a `sequence<string>` to a `std::list<std::wstring>`, Ice knows exactly what to do. However, if you select a class or template outside the standard C++ library, such as your own C++ template, you usually need to tell Ice how to marshal and unmarshal this type.

The following pages describe the set of classes and templates that Ice uses for marshaling and unmarshaling, and how to plug-in your own C++ types into this framework.

Topics

- The C++ Stream Helpers
- The cpp:type and cpp:view-type Metadata Directives

Copyright 2016, ZeroC, Inc.

The C++ Stream Helpers

Ice for C++ uses a set of C++ templates and specializations of these templates to marshal and unmarshal C++ types. A type that can be marshaled and/or unmarshaled using this framework is known as a *Streamable.*

On this page:

- Marshaling and Unmarshaling C++ Objects
- StreamableTraits
- Marshaling and Unmarshaling Optional Values

Marshaling and Unmarshaling C++ Objects

Ice for C++ marshals C++ objects into streams of bytes, and unmarshals these C++ objects from streams of bytes. Ice for C++ has an internal implementation for these streams, and a publicly available implementation, and both implementations have mostly the same API.

To marshal a C++ object into a stream, you (or the generated code) can simply call `write(obj)` on that stream. For basic types, such as `int` or `string`, there is a corresponding write function on the stream. For other types, `write` is an inline template function that delegates to a specialization of the `StreamHelper` template class:

C++

```
namespace Ice {
class OutputStream : public IceUtil::Shared {
public:
...
    template<typename T> inline void write(const T& v) {
        StreamHelper<T, StreamableTraits<T>::helper>::write(this, v);
    }
};
}
```

Likewise, when unmarshaling a C++ object from a stream, you (or the generated code) can simply call `read(obj)` on that stream, and the reading of most types delegates to a specialization of this `StreamHelper` template class:

C++

```
namespace Ice {
class InputStream : public IceUtil::Shared {
public:
...
    template<typename T> inline void read(T& v) {
        StreamHelper<T, StreamableTraits<T>::helper>::read(this, v);
   }
};
}
```

Ice looks for the `StreamHelper` specializations in namespace Ice, so they must all be in namespace Ice.

Each `StreamHelper` specialization typically provides a static `write` function to write the C++ object into the steam, and a static `read` function to read the C++ object from the stream. The stream object in these `StreamHelper` static functions is represented by a template parameter type, which allows a `StreamHelper` specialization to work with any stream implementation:

Copyright 2016, ZeroC, Inc.

C++

```
// Typical StreamHelper specialization
namespace Ice {
template<>
struct StreamHelper<SomeType, StreamHelperCategoryXXX> {
    template<class S> static inline void
    write(S* stream, const SomeType& v) {
        ... marshal v...
    }
    template<class S> static inline void
    read(S* stream, SomeType& v) {
       ... unmarshal v ...
    }
};
}
```

The base `StreamHelper` template class is not defined, since there is no default way to read a C++ object from a stream or write this object into a stream:

C++

```
namespace Ice {
template<typename T, StreamHelperCategory st> struct StreamHelper;
}
```

For most types, Ice for C++ provides or generates the corresponding `StreamHelper` specialization (or partial specialization), for example all Slice sequences mapped to a C++ vector, C++ list or similar type use this specialization:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {
template<typename T>
struct StreamHelper<T, StreamHelperCategorySequence> {
    template<class S> static inline void
    write(S* stream, const T& v) {
        stream->writeSize(static_cast<Int>(v.size()));
        for(typename T::const_iterator p = v.begin(); p != v.end(); ++p) {
            stream->write(*p);
        }
    }
    template<class S> static inline void
    read(S* stream, T& v) {
        Int sz = stream->readAndCheckSeqSize(StreamableTraits<typename
T::value_type>::minWireSize);
        T(sz).swap(v);
        for(typename T::iterator p = v.begin(); p != v.end(); ++p) {
            stream->read(*p);
        }
    }
};
}
```

StreamableTraits

The `StreamHelper` specializations use specializations of the `StremableTraits` template class to retrieve some characteristics (or traits) of the C++ type it operates on:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {

typedef int StreamHelperCategory;
const StreamHelperCategory StreamHelperCategoryUnknown = 0;
const StreamHelperCategory StreamHelperCategoryBuiltin = 1;
const StreamHelperCategory StreamHelperCategoryStruct = 2;
const StreamHelperCategory StreamHelperCategoryStructClass = 3;
const StreamHelperCategory StreamHelperCategoryEnum = 4;
const StreamHelperCategory StreamHelperCategorySequence = 5;
const StreamHelperCategory StreamHelperCategoryDictionary = 6;
const StreamHelperCategory StreamHelperCategoryProxy = 7;
const StreamHelperCategory StreamHelperCategoryClass = 8;
const StreamHelperCategory StreamHelperCategoryUserException = 9;

template<typename T, typename Enabler = void>
struct StreamableTraits {
    static const StreamHelperCategory helper = ...;
    static const int minWireSize = 1;
    static const bool fixedLength = false;
};
}
```

In particular, `StreamHelper` specializations are divided in various categories, to allow a single partial specialization of `StreamHelper` to handle several similar C++ types. For example, Ice for C++ provides a single `StreamHelper` partial specialization for all `enum` types, with category `StreamHelperCategoryEnum`:

C++

```
template<typename T>
struct StreamHelper<T, StreamHelperCategoryEnum> {
    template<class S> static inline void
    write(S* stream, const T& v) {
        ...sanity check...
        stream->writeEnum(static_cast<Int>(v), StreamableTraits<T>::maxValue);
    }
    template<class S> static inline void
    read(S* stream, T& v) {
        Int value = stream->readEnum(StreamableTraits<T>::maxValue);
        ..sanity check...
        v = static_cast<T>(value);
    }
};
```

`StreamHelperCategory` is an integer (`int`). Values between 0 and 20 are reserved for Ice; you can use other values for your own `StreamHelper` specializations.

A `StreamableTraits` is always defined in namespace Ice and provides three `static const` data members:

Copyright 2016, ZeroC, Inc.

- `helper`
  The category of `StreamHelper` associated with this C++ type (see above). The base `StreamableTraits` template class computes `helper` automatically for maps, lists, vectors and similar types.

- `minWireSize`
  The minimum size (in bytes) of the corresponding Slice type when marshaled using the Ice Encoding. The base `StreamableTraits` template class provides the default value, 1 byte. This value should be one for sequences and dictionaries, since an empty sequence or dictionary is marshaled as a single byte (size = 0).

- `fixedLength`
  `true` when the corresponding Slice type is marshaled on a fixed number of bytes, and `false` when it is marshaled on a variable number of bytes. The default, provided by the base `StreamableTraits` template class, is false.

Ice provides `StreamableTraits` specializations for all the C++ types it uses when mapping Slice types to C++. For example, Ice provides a `StreamTraits<Ice::Long>` (for 64-bit signed integers) and a `StreamableTraits< ::std::string>`:

C++

```
namespace Ice {
template<>
struct StreamableTraits<Long> {
    static const StreamHelperCategory helper = StreamHelperCategoryBuiltin;
    static const int minWireSize = 8;
    static const bool fixedLength = true;
};

template<>
struct StreamableTraits< ::std::string> {
    static const StreamHelperCategory helper = StreamHelperCategoryBuiltin;
    static const int minWireSize = 1;
    static const bool fixedLength = false;
};
}
```

For each constructed type, such as Slice `structs`, the Slice to C++ translator generates the corresponding `StreamableTraits` specialization. For example:

C++

```
// Generated C++ code
namespace Ice {
template<>
struct StreamableTraits< ::Ice::Identity> {
    static const StreamHelperCategory helper = StreamHelperCategoryStruct;
    static const int minWireSize = 2;
    static const bool fixedLength = false;
};
```

Marshaling and Unmarshaling Optional Values

As described in Data Encoding for Optional Values, the encoding or wire representation of an optional data member or parameter consists of:

Copyright 2016, ZeroC, Inc.

- An optional format (a 3-bits value) that depends on the associated Slice type. For example, the optional format for a `sequence<string>` is `FSize` (or `6`).
- The tag associated with this optional data member or parameter (this tag is a positive integer)
- The actual value of this optional data member or parameter

To marshal an optional value into a stream, you (or the generated code) call `write(tag, obj)` on that stream, where `tag` is the optional's tag (an int), and obj is the optional value. `write` is a template function that uses a specialization of the `StreamOptionalHelper` template class when this optional value is set:

C++

```
namespace Ice {
class OutputStream : public IceUtil::Shared {
public:
...
  template<typename T> inline void write(Int tag, const IceUtil::Optional<T>& v) {
        if(v) {
            writeOptional(tag, StreamOptionalHelper<T,
                                                    StreamableTraits<T>::helper,

StreamableTraits<T>::fixedLength>::optionalFormat);
            StreamOptionalHelper<T, StreamableTraits<T>::helper,
StreamableTraits<T>::fixedLength>::write(this, *v);
        }
    }
};
}
```

Likewise, when unmarshaling an optional value from a stream, you (or the generated code) call `read(tag, obj)` on that stream, which delegates to a specialization of the same `StreamOptionalHelper` template class when the stream holds a value for this optional data member or parameter:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {
class InputStream : public IceUtil::Shared {
public:
...
    template<typename T> inline void read(Int tag, IceUtil::Optional<T>& v) {
        if(readOptional(tag, StreamOptionalHelper<T,
                                                  StreamableTraits<T>::helper,

StreamableTraits<T>::fixedLength>::optionalFormat)) {
            v.__setIsSet();
            StreamOptionalHelper<T, StreamableTraits<T>::helper,
StreamableTraits<T>::fixedLength>::read(this, *v);
        }
        else {
            v = IceUtil::None;
        }
    }
};
}
```

Each `StreamOptionalHelper` specialization usually provides `static read` and `write` functions, just like `StreamHelper` specializations, and also an `optionalFormat` static data member that provides the optional format of the corresponding Slice type.

For many types, the `StreamOptionalHelper` specialization simply delegates to this type's `StreamHelper` specialization, and computes `optionalFormat` with the `GetOptionalFormat` template:

Copyright 2016, ZeroC, Inc.

C++

```
// GetOptionalFormat is declared but never defined
template<StreamHelperCategory st, int minWireSize, bool fixedLength>
struct GetOptionalFormat;

// Base StreamOptionalHelper template class
namespace Ice {
template<typename T, StreamHelperCategory st, bool fixedLength>
struct StreamOptionalHelper {
    typedef StreamableTraits<T> Traits;

    static const OptionalFormat optionalFormat = GetOptionalFormat<st,
Traits::minWireSize, fixedLength>::value;

    template<class S> static inline void
    write(S* stream, const T& v) {
        stream->write(v);
    }
    template<class S> static inline void
    read(S* stream, T& v) {
        stream->read(v);
    }
};
}
```

For example, the optional format for an Ice::Long (64-bit signed integer, encoded on exactly 8 bytes) is provided by this specialization of `GetOptionalFormat`:

C++

```
namespace Ice {
template<>
struct GetOptionalFormat<StreamHelperCategoryBuiltin, 8, true> {
    static const OptionalFormat value = OptionalFormatF8;
};
}
```

See Also

- Data Encoding
- Dynamic Ice

Copyright 2016, ZeroC, Inc.

The cpp:type and cpp:view-type Metadata Directives

Ice for C++ provides two metadata directives that allow you to map Slice types to arbitrary C++ types: `cpp:type` and `cpp:view-type`. The se metadata directives currently apply only to the following Slice types:

- `string`
- `sequence`
- `dictionary`

On this page:

- Customizing the sequence mapping with cpp:type
- Customizing the dictionary mapping with cpp:type
- cpp:type with custom C++ types
- cpp:type:string and cpp:type:wstring
- Avoiding copies with cpp:view-type
- Using both cpp:type and cpp:view-type

Customizing the sequence mapping with cpp:type

The `cpp:type:`*`c++-type`* metadata directive allows you to map a given Slice type, data member or parameter to the C++ type of your choice.

For example, you can override the default mapping of a Slice sequence type:

Slice

```
[["cpp:include:list"]]

module Food {
    enum Fruit { Apple, Pear, Orange };

    ["cpp:type:std::list< ::Food::Fruit>"]
    sequence<Fruit> FruitPlatter;
};
```

With this metadata directive, the Slice sequence now maps to a C++ `std::list` instead of the default `std::vector`:

C++

```
#include <list>

namespace Food {

    typedef std::list< Food::Fruit> FruitPlatter;

    // ...
}
```

The Slice to C++ compiler takes the string following the `cpp:type:` prefix as the name of the mapped C++ type. For example, we could use `["cpp:type:::std::list< ::Food::Fruit>"]`. In that case, the compiler would use a fully-qualified name to define the type:

Copyright 2016, ZeroC, Inc.

C++

```
typedef ::std::list< ::Food::Fruit> FruitPlatter;
```

Note that the code generator inserts whatever string you specify following the `cpp:type:` prefix literally into the generated code. We recommend you use fully qualified names to avoid C++ compilation failures due to unknown symbols.

Also note that, to avoid compilation errors in the generated code, you must instruct the compiler to generate an appropriate include directive with the `cpp:include` global metadata directive. This causes the compiler to add the line

C++

```
#include <list>
```

to the generated header file.

In addition to modifying the type of a sequence itself, you can also modify the mapping for particular return values or parameters. For example:

Slice

```
[["cpp:include:list"]]
[["cpp:include:deque"]]

module Food {
    enum Fruit { Apple, Pear, Orange };

    sequence<Fruit> FruitPlatter;

    interface Market {
        ["cpp:type:list< ::Food::Fruit>"]
        FruitPlatter barter(["cpp:type:deque< ::Food::Fruit>"] FruitPlatter offer);
    };
};
```

With this definition, the default mapping of `FruitPlatter` to a C++ `vector` still applies but the return value of `barter` is mapped as a `list`, and the `offer` parameter is mapped as a `deque`.

Instead of `std::list` or `std::deque`, you can specify a type of your own as the sequence type, for example:

Copyright 2016, ZeroC, Inc.

Slice

```
[["cpp:include:FruitBowl.h"]]

module Food {
    enum Fruit { Apple, Pear, Orange };

    ["cpp:type:FruitBowl"]
    sequence<Fruit> FruitPlatter;
};
```

With these metadata directives, the compiler will use a C++ type `FruitBowl` as the sequence type, and add an `include` directive for the header file `FruitBowl.h` to the generated code.

The class or template class you provide must meet the following requirements:

- The class must have a default constructor.
- The class must have a copy constructor.

If you use a class that also meets the following requirements

- The class has a single-argument constructor that takes the size of the sequence as an argument of unsigned integral type.
- The class has a member function `size` that returns the number of elements in the sequence as an unsigned integral type.
- The class provides a member function `swap` that swaps the contents of the sequence with another sequence of the same type.
- The class defines `iterator` and `const_iterator` types and provides `begin` and `end` member functions with the usual semantics; its iterators are comparable for equality and inequality.

then you do not need to provide code to marshal and unmarshal your custom sequence – Ice will do it automatically.

Less formally, this means that if the provided class looks like a `vector`, `list`, or `deque` with respect to these points, you can use it as a custom sequence implementation without any additional coding.

Customizing the dictionary mapping with cpp:type

You can override the default mapping of Slice dictionaries to C++ maps with a `cpp:type` metadata directive, for example:

Slice

```
[["cpp:include:unordered_map"]]

["cpp:type:std::unordered_map<Ice::Long, Employee>"] dictionary<long, Employee>
EmployeeMap;
```

With this metadata directive, the dictionary now maps to a C++ `std::unordered_map`:

C++

```
#include <unordered_map>

typedef std::unordered_map<Ice::Long, Employee> EmployeeMap;
```

Like with sequences, anything following the `cpp:type:` prefix is taken to be the name of the type. For example, we could use `["cpp:type:::std::unordered_map<int, std::string>"]`. In that case, the compiler would use a fully-qualified name to define the type:

Copyright 2016, ZeroC, Inc.

C++

```
typedef ::std::unordered_map<int, std::string> IntStringDict;
```

To avoid compilation errors in the generated code, you must instruct the compiler to generate an appropriate include directive with the `cpp:include` global metadata directive. This causes the compiler to add the line

C++

```
#include <unordered_map>
```

to the generated header file.

Instead of `std::unordered_map`, you can specify a type of your own as the dictionary type, for example:

Slice

```
[["cpp:include:CustomMap.h"]]

["cpp:type:MyCustomMap<Ice::Long, Employee>"] dictionary<long, Employee>
EmployeeMap;
```

With these metadata directives, the compiler will use a C++ type `MyCustomMap` as the dictionary type, and add an include directive for the header file `CustomMap.h` to the generated code.

The class or template class you provide must meet the following requirements:

- The class must have a default constructor.
- The class must have a copy constructor.
- The class must provide nested types named `key_type`, `mapped_type` and `value_type`.
- The class must provide `iterator` and `const_iterator` types and provide `begin` and `end` member functions with the usual semantics; these iterators must be comparable for equality and inequality.
- The class must provide a `clear` function.
- The class must provide an `insert` function that takes an `iterator` (as location hint) plus a `value_type` parameter, and returns an `iterator` to the new entry or to the existing entry with the given key.

Less formally, this means you can use any class or template class that looks like a standard `map` or `unordered_map` as your custom dictionary type.

In addition to modifying the type of a dictionary itself, you can also modify the mapping for particular return values or parameters. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
[["cpp:include:unordered_map"]]

module HR {

    struct Employee {
       long   number;
       string firstName;
       string lastName;
    };
    dictionary<long, Employee> EmployeeMap;

    interface Office {
        ["cpp:type:std::unordered_map<Ice::Long, Employee>"] EmployeeMap
getAllEmployees();
    };
};
```

With this definition, `getAllEmployees` returns an `unordered_map`, while other unqualified parameters of type `EmployeeMap` would use the default mapping (to a `std::map`).

cpp:type with custom C++ types

If your C++ type does not look like a standard C++ container, you need to tell Ice how to marshal and unmarshal this type by providing your own `StreamHelper` specialization for this type.

For example, you can map a Slice `sequence<byte>` to a Google Protocol Buffer C++ class, `tutorial:Person`, with the following metadata directive:

Slice

```
module Demo {
   ["cpp:type:tutorial::Person"] sequence<byte> Person;
};
```

Since `tutorial::Person` does not look like a `vector<Ice::Byte>` or a `list<Ice::Byte>`, you need a `StreamHelper` specialization for `tutorial::Person`.

The simplest is to create a `StreamHelper` specialization that handles only this specific class:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {

template<>
struct StreamHelper<tutorial::Person, StreamHelperCategoryUnknown> {
    template<class S> static inline void
    write(S* stream, const tutorial::Person& v) {
        // ... marshal v into a sequence of bytes...
    }

    template<class S> static inline void
    read(S* stream, tutorial::Person& v) {
        //... unmarshal bytes from stream into v...
    }
};
}
```

You should also provide the corresponding `StreamTraits` specialization:

C++

```
namespace Ice {
template<>
struct StreamableTraits<tutorial::Person> {
    static const StreamHelperCategory helper = StreamHelperCategoryUnknown;
    static const int minWireSize = 1;
    static const bool fixedLength = false;
};
}
```

This `StreamTraits` specialization is actually optional for `tutorial::Person`, and more generally for any mapping of `sequence<byte>`, since these are the values provided by the base `StreamableTraits` template.

Finally, remember to insert the header for these `StreamHelper` and `StreamTraits` specializations with a `cpp:include` global metadata directive:

Slice

```
[["cpp:include:Person.pb.h"]]
[["cpp:include:PersonStreaming.h"]]
module Demo {
["cpp:type:tutorial::Person"] sequence<byte> Person;
```

Now, if your application maps several Slice `sequence<byte>` to different Google Protocol Buffers, you could provide a `StreamHelper` specialization for each of these Google Protocol Buffer classes, but it would be more judicious to provide a single partial specialization capable of marshaling and unmarshaling any Google Protocol Buffer C++ class as a Slice `sequence<byte>` in a stream:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {

// A new helper category for all Google Protocol Buffers
const StreamHelperCategory StreamHelperCategoryProtobuf = 100;

// All classes derived from ::google::protobuf::MessageLite will use this
StreamableTraits
template<typename T>
struct StreamableTraits<T, typename std::enable_if<std::is_base_of<
::google::protobuf::MessageLite, T>::value >::type> {
    static const StreamHelperCategory helper = StreamHelperCategoryProtobuf;
    static const int minWireSize = 1;
    static const bool fixedLength = false;
};
// T can be any Google Protocol Buffer C++ class
template<typename T>
struct StreamHelper<T, StreamHelperCategoryProtobuf> {
    template<class S> static inline void
    write(S* stream, const T& v) {
        std::vector<Byte> data(v.ByteSize());
        // ... marshal v into a sequence of bytes...
        stream->write(&data[0], &data[0] + data.size());
    }

    template<class S> static inline void
    read(S* stream, T& v) {
        std::pair<const Byte*, const Byte*> data;
        stream->read(data);
        //... unmarshal data into v...
    }
};
}
```

If you use any of these `sequence<byte>` mapped to Google Protocol Buffers for optional data members or optional parameters, you also need to tell Ice which optional format to use:

C++

```
namespace Ice {

// Optional format for Slice sequence<byte> mapped to Google Protocol Buffer
// The template parameters correspond to the data members of the corresponding
StreamTraits specialization.
template<>
struct GetOptionalFormat<StreamHelperCategoryProtobuf, 1, false> {
    static const OptionalFormat value = OptionalFormatVSize;
};
}
```


Copyright 2016, ZeroC, Inc.

The `OptionalFormat` provided by `GetOptionalFormat` for `StreamHelperCategoryUnknown` and its default `StreamableTraits` is `OptionalFormatVSize`, like in the example above. This way, if your application needs a single `sequence<byte>` mapped to a custom C++ type, you can provide just a `StreamHelper` specialization for this type and rely on the default `StreamTraits` and `GetOptionalFormat` templates.

cpp:type:string and cpp:type:wstring

The metadata directives `cpp:type:string` and `cpp:type:wstring` are used map Slice `strings` to `std::string` or `std::wstring`, as described in Alternative String Mapping for C++. The Slice to C++ compiler recognize `string` and `wstring` as special tokens and does not treat them like C++ types with these names.

For example, you can map a `sequence<string>` to a `std::vector<std::wstring>` with either of the following directives:

Slice

```
// Special cpp:type:wstring metadata directive: it maps the sequence to a
std::vector<std::wstring>, not to a wstring!
["cpp:type:wstring"] sequence<string>;
```

or

Slice

```
// Maps the sequence to a std::vector<std::wstring> using a regular cpp:type
metadata directive
["cpp:type:std::vector<std::wstring>"] sequence<string>;
```

Avoiding copies with cpp:view-type

The main drawback of using standard C++ library classes to represent strings, sequences and dictionaries (as Ice does by default) is copying. For example, if we transfer an array of characters with Ice:

Slice

```
void sendChars(string s);
```

with the default mapping, our C++ code would look like:

C++

```
// client side
const char* cstring = ... null terminated array of chars;
proxy->sendChars(string(cstring));

// server side
void sendChars(const string& s, ...) {
   // use s;
}
```

Each invocation triggers a number of copies:

Copyright 2016, ZeroC, Inc.

- (client) the creation of the string on the client side makes a copy of the array of characters
- (client) `sendChars` copies the characters of the string into the client-side marshaling buffer
- (server) the unmarshaling code creates a string from the bytes in the server-side marshaling buffer

If instead of `std::string`, Ice could use a string-type with view semantics–its instances point to memory owned by some other objects–we could avoid these copies. This is exactly what the `cpp:view-type` metadata directive allows us to do: select a custom mapping for Slice `strings`, `sequences` and `dictionaries`.

For example, we can map our Slice string parameter in the example above to a `std::experimental::string_view`, that has such view semantics:

Slice

```
void sendChars(["cpp:type:std::experimental::string_view"] string s);
```

Our C++ code pretty much the same, but we avoid several copies:

C++

```
// client side
const char* cstring = ... null terminated array of chars;
proxy->sendChars(string_view(cstring)); // string_view points to the characters
cstring, without copy

// server side
void sendChars(const string_view& s, ...) { // string_view points to bytes in the
server-side unmarshaling buffer
   // use s;
}
```

With this metadata directive, the only remaining copy during each invocation is:

- (client) `sendChars` copies the characters into the client-side marshaling buffer

This `cpp:view-type` metadata is very much like the `cpp:type` metadata directive described earlier, and just like for `cpp:type`, the Slice to C++ compiler uses the string provided after the `cpp:view-type:` prefix literally in the generated code. The compiler does not (and cannot) check that this string represents a valid C++ type, or a C++ type with view semantics.

There are however two significant differences between `cpp:view-type` and `cpp:type:`

- `cpp:view-type` can be applied only to operation parameters, while `cpp:type` can be applied to the definition of sequence and dictionary types, to operation parameters and to data members
- `cpp:view-type` changes the mapping of a parameter to the specified C++ type only when it is safe to use a view object (an object that does not own memory), while `cpp:type` changes the mapping all the time.

For example, if instead of sending an array of characters from a client to a server, we return an array of character from the server to the client, without using AMI or AMD (the default):

Slice

```
string getChars();
```

With the default mapping, we get a `std::string` back, and each invocation makes a few copies:

Copyright 2016, ZeroC, Inc.

C++

```
// client-side
string s = prx->getChars();

// server side
string getChars(...) {
   const char* cstring = ... null terminated array of chars;
   return cstring; // the conversion to string makes a copy
}
```

Now, if we map the returned string to a `string_view` with the `cpp:type` metadata directive:

Slice

```
["cpp:type:std::experimental::string_view"] string getChars();
```

we avoid some copies but our `string_view` now points to deallocated memory and our program will crash!

C++

```
// client-side
string_view s = prx->getChars(); // string_view points to the client-side
marshaling buffer, which is deallocated as soon as getChars() returns

// server side
string_view getChars(...) {
   const char* cstring = ... null terminated array of chars;
   return cstring; // string_view points to (or may point to) stack-allocated
memory, reclaimed when getChars() returns
}
```


Never use the `cpp:type` metadata directive with a view type (a type that does not manage its memory); you should only use `cpp:view-type` with view types.

With `cpp:view-type`, the compiler changes the mapping only when it's safe to use a view-type, namely for:

- Input parameters, on the client-side and on the server-side
- Out and return parameters provided by the Ice run-time to AMI type-safe callbacks and AMI lambdas
- Out and return parameters provided to AMD callbacks

The `cpp:array` and `cpp:range` metadata directives for sequences follow the same rules.

With our `getChars` example:

Slice

```
["cpp:view-type:std::experimental::string_view"] string getChars();
```


Copyright 2016, ZeroC, Inc.

it is not safe the change the client-side mapping or the server-side mapping (unless we use AMD), so the returned C++ type remains a `std::string`.

Like with the `cpp:type` metadata directive, if the C++ type specified with `cpp:view-type` is a standard library type (such as `std::list`) or looks like one, Ice will automatically marshal and unmarshal it for you. You just need to include the corresponding header with the `cpp:include` global metadata directive:

Slice

```
[["cpp:include:MyContainer.h"]]
module Sample {
    ...
};
```

For other C++ types, you need to tell Ice how to marshal and unmarshal these types, as described above in cpp:type with custom C++ types.

In particular, Ice does not know how to marshal and marshal the experimental `string_view` type. If you want to map some string parameters to `string_view`, you need to provide a `StreamHelper` for `string_view`, such as:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {

template<>
struct StreamableTraits<std::experimental::string_view> {
    static const StreamHelperCategory helper = StreamHelperCategoryBuiltin;
    static const int minWireSize = 1;
    static const bool fixedLength = false;
};

template<>
struct StreamHelper<std::experimental::string_view, StreamHelperCategoryBuiltin> {
    template<class S> static inline void
    write(S* stream, const std::experimental::string_view& v) {
#ifdef STRING_VIEW_IGNORE_STRING_CONVERTER
        stream->write(v.data(), v.size(), false);
#else
        stream->write(v.data(), v.size(), true);
#endif
    }
    template<class S> static inline void
    read(S* stream, std::experimental::string_view& v) {
        const char* vdata = 0;
        size_t vsize = 0;
#ifdef STRING_VIEW_IGNORE_STRING_CONVERTER
        stream->read(vdata, vsize);
#else
        std::string holder;
        stream->read(vdata, vsize, holder);

        // If holder is not empty, a string conversion occured, and we can't return
a
        // string_view since it does not hold the memory
        if(!holder.empty()) {
            throw Ice::MarshalException(__FILE__, __LINE__, "string conversion not
supported");
        }
#endif
        if(vsize > 0) {
            v = std::experimental::string_view(vdata, vsize);
        }
        else {
            v.clear();
        }
    }
};
}
```

This `StreamHelper` specialization will take care of plain `string` parameters, and also sequence or dictionary parameters that contain `strings` mapped to `string_view` (when it's safe to do so), such as:

Copyright 2016, ZeroC, Inc.

Slice

```
["amd", "cpp:view-type:std::vector<std::experimental::string_view>"] StringSeq
 echoStringSeq(["cpp:view-type:std::vector<std::experimental::string_view>"]
StringSeq seq);
```

The Ice/throughput demo illustrates the use of `cpp:view-type` with sequences of strings.

Using both cpp:type and cpp:view-type

If you specify both `cpp:type` and `cpp:view-type` for the same parameter, `cpp:view-type` applies to mapped parameters safe for view types, and `cpp:type` applies to all other mapped parameters.

With the following somewhat contrived example:

Slice

```
["amd", "cpp:view-type:std::vector<std::experimental::string_view>",
"cpp:type:std::list<std::wstring>"] StringSeq getStringSeq();
```

calling `getStringSeq()` on a proxy returns a `std::list<std::wstring>`, while the `StringSeq` is passed to the AMD callback as a `std::vector<string_view>`.

See Also

- Data Encoding
- Dynamic Ice

Copyright 2016, ZeroC, Inc.

Version Information in C++

The header file `IceUtil/Config.h` defines two macros that expand to the version of the Ice run time:

C++

```
#define ICE_STRING_VERSION "3.6.2" // "<major>.<minor>.<patch>"
#define ICE_INT_VERSION 30602      // AABBCC, with AA=major,
                                   // BB=minor, CC=patch
```

`ICE_STRING_VERSION` is a string literal in the form <*major*>.<*minor*>.<*patch*>, for example, `3.6.2`. For beta releases, the version is <*major*>.<*minor*>b, for example, `3.6b`.

`INT_VERSION` is an integer literal in the form *AABBCC*, where *AA* is the major version number, *BB* is the minor version number, and *CC* is the patch level, for example, `30602` for version 3.6.2. For beta releases, the patch level is set to 51 so, for example, for version 3.6b, the value is `30651`.

Copyright 2016, ZeroC, Inc.

slice2cpp Command-Line Options

On this page:

- slice2cpp Command-Line Options
  - --header-ext EXT
  - --source-ext EXT
  - --add-header HDR[,GUARD]
  - --include-dir DIR
  - --impl
  - --dll-export MACRO
  - --checksum
  - --stream
- Include Directives
  - Header Files
  - Source Files

`slice2cpp` Command-Line Options

The Slice-to-C++ compiler, `slice2cpp`, offers the following command-line options in addition to the standard options.

`--header-ext` *EXT*

Changes the file extension for the generated header files from the default h to the extension specified by *EXT*.

You can also change the header file extension with a global metadata directive:

Slice

```
[["cpp:header-ext:hpp"]]

// ...
```

Only one such directive can appear in each source file. If you specify a header extension on both the command line and with a metadata directive, the metadata directive takes precedence. This ensures that included Slice files that were compiled separately get the correct header extension (provided that the included Slice files contain a corresponding metadata directive). For example:

Slice

```
// File example.ice
#include <Ice/BuiltinSequences.ice>

// ...
```

Compiling this file with

```
$ slice2cpp --header-ext=hpp -I/opt/Ice/include example.ice
```

generates `example.hpp`, but the `#include` directive in that file is for `Ice/BuiltinSequences.h` (not `Ice/BuiltinSequences.hpp`) because `BuiltinSequences.ice` contains the metadata directive `[["cpp:header-ext:h"]]`.

You normally will not need to use this metadata directive. The directive is necessary only if:

- You `#include` a Slice file in one of your own Slice files.
- The included Slice file is part of a library you link against.
- The library ships with the included Slice file's header.

Copyright 2016, ZeroC, Inc.

- The library header uses a different header extension than your own code.

For example, if the library uses `.hpp` as the header extension, but your own code uses `.h`, the library's Slice file should contain a `[["cpp:header?ext:hpp"]]` directive. (If the directive is missing, you can add it to the library's Slice file.)

`--source-ext` *`EXT`*

Changes the file extension for the generated source files from the default `cpp` to the extension specified by *`EXT`*.

`--add-header` *`HDR`*`[,`*`GUARD`*`]`

This option adds an include directive for the specified header at the beginning of the generated source file (preceding any other include directives). If *`GUARD`* is specified, the include directive is protected by the specified guard. For example, `--add-header precompiled.h,__PRECOMPILED_H__` results in the following directives at the beginning of the generated source file:

C++

```
#ifndef __PRECOMPILED_H__
#define __PRECOMPILED_H__
#include <precompiled.h>
#endif
```

The option can be repeated to create include directives for several files.

As suggested by the preceding example, this option is useful mainly to integrate the generated code with a compiler's precompiled header mechanism.

`--include-dir` *`DIR`*

Modifies `#include` directives in source files to prepend the path name of each header file with the directory *`DIR`*.

`--impl`

Generate sample implementation files. This option will not overwrite an existing file.

`--dll-export` *`MACRO`*

Use *`MACRO`* to control the export and import of symbols from DLLs on Windows and dynamic shared libraries on other platforms. This option allows you to export symbols from the generated code, and place such generated code in a DLL (on Windows) or shared library (on other platforms). As an example, compiling a Slice definition with:

```
$ slice2cpp --dll-export WIDGET_API x.ice
```

results in the following additional code being generated into `x.h`:

C++

```
#ifndef WIDGET_API
#   ifdef WIDGET_API_EXPORTS
#       define WIDGET_API ICE_DECLSPEC_EXPORT
#   else
#       define WIDGET_API ICE_DECLSPEC_IMPORT
#   endif
#endif
```


Copyright 2016, ZeroC, Inc.

The generated code also includes the provided `MACRO` name (`WIDGET_API` in our example) in the declaration of classes and functions that need to be exported (when building a DLL or dynamic library) or imported (when using such library).

`ICE_DECLSPEC_EXPORT` and `ICE_DECLSPEC_IMPORT` are macros that expand to compiler-specific attributes. For example, for Visual Studio, they are defined as:

C++

```
#if defined(_WIN32) && !defined(ICE_STATIC_LIBS)
#   define ICE_DECLSPEC_EXPORT __declspec(dllexport)
#   define ICE_DECLSPEC_IMPORT __declspec(dllimport)
...
```

With GCC and clang, they are defined as:

C++

```
#elif defined(__GNUC__)
#   define ICE_DECLSPEC_EXPORT __attribute__((visibility ("default")))
#   define ICE_DECLSPEC_IMPORT __attribute__((visibility ("default")))
...
```

The net effect is that, if you are using Visual Studio to create a DLL that includes `x.cpp`, and you want to be able to use `x.cpp`'s definitions from code outside this DLL, you need to export `x.cpp`'s symbols by compiling `x.cpp` with `-DWIDGET_API_EXPORTS`. If you use GCC or clang, you can likewise compile `x.cpp` with `-DWIDGET_API_EXPORTS`; this is however optional since GCC and clang's definitions for `ICE_DECLSPEC_IMPORT` and `ICE_DECLSPEC_EXPORT` are identical.

Similar definitions exist for other platforms. For platforms that do not have any concept of explicit export or import of shared library symbols, both macros are empty.

`--checksum`

Generate checksums for Slice definitions.

`--stream`

Generate streaming helper functions for Slice classes and exceptions. For other Slice types, streaming support is always generated.

Include Directives

The `#include` directives generated by the Slice-to-C++ compiler can be a source of confusion if the semantics governing their generation are not well-understood. The generation of `#include` directives is influenced by the command-line options `-I` and `--include-dir`; these options are discussed in more detail below. The `--output-dir` option directs the translator to place all generated files in a particular directory, but has no impact on the contents of the generated code.

Given that the `#include` directives in header files and source files are generated using different semantics, we describe them in separate sections.

Header Files

In most cases, the compiler generates the appropriate `#include` directives by default. As an example, suppose file `A.ice` includes `B.ice` using the following statement:

Copyright 2016, ZeroC, Inc.

Slice

```
// A.ice
#include <B.ice>
```

Assuming both files are in the current working directory, we run the compiler as shown below:

```
$ slice2cpp -I. A.ice
```

The generated file `A.h` contains this `#include` directive:

C++

```
// A.h
#include <B.h>
```

If the proper include paths are specified to the C++ compiler, everything should compile correctly.

Similarly, consider the common case where `A.ice` includes `B.ice` from a subdirectory:

Slice

```
// A.ice
#include <inc/B.ice>
```

Assuming both files are in the `inc` subdirectory, we run the compiler as shown below:

```
$ slice2cpp -I. inc/A.ice
```

The default output of the compiler produces this `#include` directive in `A.h`:

C++

```
// A.h
#include <inc/B.h>
```

Again, it is the user's responsibility to ensure that the C++ compiler is configured to find `inc/B.h` during compilation.

Now let us consider a more complex example, in which we do not want the `#include` directive in the header file to match that of the Slice file. This can be necessary when the organizational structure of the Slice files does not match the application's C++ code. In such a case, the user may need to relocate the generated files from the directory in which they were created, and the `#include` directives must be aligned with the new structure.

For example, let us assume that `B.ice` is located in the subdirectory `slice/inc`:

Copyright 2016, ZeroC, Inc.

Slice

```
// A.ice
#include <slice/inc/B.ice>
```

However, we do not want the `slice` subdirectory to appear in the `#include` directive generated in the header file, therefore we specify the additional compiler option `-Islice`:

```
$ slice2cpp -I. -Islice slice/inc/A.ice
```

The generated code demonstrates the impact of this extra option:

C++

```
// A.h
#include <inc/B.h>
```

As you can see, the `#include` directives generated in header files are affected by the include paths that you specify when running the compiler. Specifically, the include paths are used to abbreviate the path name in generated `#include` directives.

When translating an `#include` directive from a Slice file to a header file, the compiler compares each of the include paths against the path of the included file. If an include path matches the leading portion of the included file, the compiler removes that leading portion when generating the `#include` directive in the header file. If more than one include path matches, the compiler selects the one that results in the shortest path for the included file.

For example, suppose we had used the following options when compiling `A.ice`:

```
$ slice2cpp -I. -Islice -Islice/inc slice/inc/A.ice
```

In this case, the compiler compares all of the include paths against the included file `slice/inc/B.ice` and generates the following directive:

C++

```
// A.h
#include <B.h>
```

The option `-Islice/inc` produces the shortest result, therefore the default path for the included file (`slice/inc/B.h`) is replaced with `B.h`.

In general, the `-I` option plays two roles: it enables the preprocessor to locate included Slice files, and it provides you with a certain amount of control over the generated `#include` directives. In the last example above, the preprocessor locates `slice/inc/B.ice` using the include path specified by the `-I.` option. The remaining `-I` options do not help the preprocessor locate included files; they are simply hints to the compiler.

Finally, we recommend using caution when specifying include paths. If the preprocessor is able to locate an included file via multiple include paths, it always uses the first include path that successfully locates the file. If you intend to modify the generated `#include` directives by specifying extra `-I` options, you must ensure that your include path hints match the include path selected by the preprocessor to locate the included file. As a general rule, you should avoid specifying include paths that enable the preprocessor to locate a file in multiple ways.

Copyright 2016, ZeroC, Inc.

Source Files

By default, the compiler generates `#include` directives in source files using only the base name of the included file. This behavior is usually appropriate when the source file and header file reside in the same directory.

For example, suppose `A.ice` includes `B.ice` from a subdirectory, as shown in the following snippet of `A.ice`:

Slice

```
// A.ice
#include <inc/B.ice>
```

We generate the source file using this command:

```
$ slice2cpp -I. inc/A.ice
```

Upon examination, we see that the source file contains the following `#include` directive:

C++

```
// A.cpp
#include <B.h>
```

However, suppose that we wish to enforce a particular standard for generated `#include` directives so that they are compatible with our C++ compiler's existing include path settings. In this case, we use the `--include-dir` option to modify the generated code. For example, consider the compiler command shown below:

```
$ slice2cpp --include-dir src -I. inc/A.ice
```

The source file now contains the following `#include` directive:

C++

```
// A.cpp
#include <src/B.h>
```

Any leading path in the included file is discarded as usual, and the value of the `--include-dir` option is prepended.

See Also

- Using the Slice Compilers
- Using Slice Checksums in C++
- Streaming Interfaces

Copyright 2016, ZeroC, Inc.

C++ Strings and Character Encoding

On the wire, Ice transmits all strings as Unicode strings in UTF-8 encoding. For languages other than C++, Ice uses strings in their language-native Unicode representation and converts automatically to and from UTF-8 for transmission, so applications can transparently use characters from non-English alphabets.

However, for C++, how strings are represented inside a process depends on the platform as well as the mapping that is chosen for a particular string: the default mapping to `std::string`, or the alternative mapping to `std::wstring`.

This discussion is only relevant for C++. For scripting language mappings based on Ice for C++, it is possible to use Ice's default string converter plug-in and to install your own string converter plug-in.

We will explore how strings are encoded by the Ice for C++ run time, and how you can achieve automatic conversion of strings in their native representation to and from UTF-8. For an example of using string converters in C++, refer to the sample program provided in the `demo/Ice/converter` subdirectory of your Ice distribution.

By default, the Ice run time encodes strings as follows:

- Narrow strings (that is, strings mapped to `std::string`) are presented to the application in UTF-8 encoding and, similarly, the application is expected to provide narrow strings in UTF-8 encoding to the Ice run time for transmission.

  With this default behavior, the application code is responsible for converting between the native codeset for 8-bit characters and UTF-8. For example, if the native codeset is ISO Latin-1, the application is responsible for converting between UTF-8 and narrow (8-bit) characters in ISO Latin-1 encoding.

  Also note that the default behavior does not require the application to do anything if it only uses characters in the ASCII range. (This is because a string containing only characters in the (7-bit) ASCII range is also a valid UTF-8 string.)

- Wide strings (that is, strings mapped to `std::wstring`) are automatically encoded as Unicode by the Ice run time as appropriate for the platform. For example, for Windows, the Ice run time converts between UTF-8 and UTF-16 in little-endian representation whereas, for Linux, the Ice run time converts between UTF-8 and UTF-32 in the endian-ness appropriate for the host CPU.

  With this default behavior, wide strings are transparently converted between their on-the-wire representation and their native C++ representation as appropriate, so application code need not do anything special. (The exception is if an application uses a non-Unicode encoding, such as Shift-JIS, as its native `wstring` codeset.)

Topics

- Installing String Converters
- UTF-8 Conversion
- String Parameters in Local Calls
- Built-in String Converters
- String Conversion Convenience Functions
- The iconv String Converter
- The Ice String Converter Plug-in
- Custom String Converter Plug-ins

See Also

- The Ice Protocol
- C++ Mapping for Built-In Types

Copyright 2016, ZeroC, Inc.

Installing String Converters

The default behavior of the run time can be changed by providing application-specific string converters. If you install such converters, all Slice strings will be passed to the appropriate converter when they are marshaled and unmarshaled. Therefore, the string converters allow you to convert all strings transparently into their native representation without having to insert explicit conversion calls whenever a string crosses a Slice interface boundary.

You can install string converters by calling `IceUtil::setProcessStringConverter` for the narrow string converter, and `IceUtil::setProcessWstringConverter` for the wide string converter. Any strings that use the default (`std::string`) mapping are passed through the specified narrow string converter and any strings that use the wide (`std::wstring`) mapping are passed through the specified wide string converter.

You can also retrieve the previously installed string converters (or default string converters) with `IceUtil::getProcessStringConverter` and `IceUtil::getProcessWstringConverter`. The default narrow string converter is null, meaning all `std::string`s use the UTF-8 encoding.

The string converters are defined as follows:

C++

```
namespace IceUtil {

class UTF8Buffer {
public:
    virtual Byte* getMoreBytes(size_t howMany, Byte* firstUnused) = 0;
    virtual ~UTF8Buffer() {}
};

template<typename charT>
class BasicStringConverter : public IceUtil::Shared {
public:
    virtual Byte* toUTF8(const charT* sourceStart, const charT* sourceEnd,
                         UTF8Buffer&) const = 0;

    virtual void fromUTF8(const Byte* sourceStart,
                          const Byte* sourceEnd,
                          std::basic_string<charT>& target) const;
};

typedef BasicStringConverter<char> StringConverter;
typedef IceUtil::Handle<StringConverter> StringConverterPtr;

typedef BasicStringConverter<wchar_t> WstringConverter;
typedef IceUtil::Handle<WstringConverter> WstringConverterPtr;

StringConverterPtr getProcessStringConverter();
void setProcessStringConverter(const StringConverterPtr&);

WstringConverterPtr getProcessWstringConverter();
void setProcessWstringConverter(const WstringConverterPtr&);
}
```

As you can see, both narrow and wide string converters are simply templates with either a narrow or a wide character (`char` or `wchar_t`) as the template parameter.

Each communicator caches the narrow string converter and wide string converter installed when this communicator is initialized.

Copyright 2016, ZeroC, Inc.

You should always install your string converters before creating your communicator(s). When using a plugin to set your string converters, you need to set the string converters in the constructor of your plugin class (which is executed when the plugin is loaded) and not in the initialization function of the plugin class (which is executed after the communicator has read and cached the process-wide string converters).

See Also

- Communicator Initialization
- C++ Mapping for Built-In Types

Copyright 2016, ZeroC, Inc.

UTF-8 Conversion

On this page:

- Converting to UTF-8
- Converting from UTF-8

Converting to UTF-8

If you have installed a string converter, the Ice run time calls the converter's `toUTF8` function whenever it needs to convert a native string into UTF-8 representation for transmission. The `sourceStart` and `sourceEnd` pointers point at the first byte and one-beyond-the-last byte of the source string, respectively. The implementation of `toUTF8` must return a pointer to the first unused byte following the converted string.

Your implementation of `toUTF8` must allocate the returned string by calling the `getMoreBytes` member function of the `UTF8Buffer` class that is passed as the third argument. (`getMoreBytes` throws a `std::bad_alloc` if it cannot allocate enough memory.) The `firstUnused` parameter must point at the first unused byte of the allocated memory region. You can make several calls to `getMoreBytes` to incrementally allocate memory for the converted string. If you do, `getMoreBytes` may relocate the buffer in memory. (If it does, it copies the part of the string that was converted so far into the new memory region.) The function returns a pointer to the first unused byte of the (possibly relocated) memory.

Conversion can also fail because the encoding of the source string is internally incorrect. In that case, you should throw a `IceUtil::IllegalConversionException` exception from `toUTF8`, for example:

C++

```
throw IceUtil::IllegalConversionException(__FILE__, __LINE__, "bad encoding because
...");
```

After it has marshaled the returned string into an internal marshaling buffer, the Ice run time deallocates the string.

Converting from UTF-8

During unmarshaling, the Ice run time calls the `fromUTF8` member function on the corresponding string converter. The function converts a UTF-8 byte sequence into its native form as a `std::string` or `std::wstring`. The string into which the function must place the converted characters is passed to `fromUTF8` as the `target` parameter.

See Also

- Installing String Converters

Copyright 2016, ZeroC, Inc.

String Parameters in Local Calls

In C++, and indirectly in Python, Ruby, and PHP, all Ice local APIs are narrow-string based, meaning you could not for example recompile `Properties.ice` to get property names and values as wide strings.

Installing a narrow-string converter could cause trouble for these local calls if UTF-8 conversion occurs in the underlying implementation. For example, the `stringToIdentity` operation creates an intermediary UTF-8 string. If this string contains characters that are not in your native codeset (as determined by the narrow-string converter), the `stringToIdentity` call will fail.

Likewise, when Ice reads properties from a configuration file, it converts the input (UTF-8 characters) into native strings. This conversion can also fail if the native encoding cannot convert some characters.

Most strings in local calls are never problematic because Ice does not perform any conversion, for example:

- adapter names in `createObjectAdapter`
- property names and values in `Properties`
- `ObjectAdapter::createProxy`, where the identity conversion occurs only when the proxy is marshaled

Finally, consider the Slice type `Ice::Context`, which is mapped in C++ as a `map<string, string>`. The mapping for `Context` cannot be changed to `map<wstring, wstring>`, therefore you cannot send or receive any context entry that is not in your narrow-string native encoding when a narrow-string converter is installed.

See Also

- Object Identity
- Properties and Configuration
- Request Contexts

Copyright 2016, ZeroC, Inc.

Built-in String Converters

Ice provides three string converters to cover common conversion requirements:

- `UnicodeWstringConverter`
  This is a string converter that converts between Unicode wide strings and UTF-8 strings. Unless you install a different string converter, this is the default converter that is used for wide strings.
- `IconvStringConverter` (Linux and Unix only)
  The `iconv` string converter converts strings using the Linux and Unix `iconv` conversion facility. It can be used to convert either wide or narrow strings.
- `WindowsStringConverter` (Windows only)
  This string converter converts between multi-byte and UTF-8 strings and uses `MultiByteToWideChar` and `WideCharToMultiByte` for its implementation.

These string converters are defined in the `IceUtil` namespace.

See Also

- The iconv String Converter

Copyright 2016, ZeroC, Inc.

String Conversion Convenience Functions

The `IceUtil` namespace contains two helper functions that allow you to convert between narrow strings and UTF-8 encoded strings using a string converter:

C++

```
std::string nativeToUTF8(const std::string&, const StringConverterPtr&);

std::string UTF8ToNative(const std::string&, const StringConverterPtr&);
```

No conversion is performed when the provided string converter is null.

The `IceUtil` namespace contains two additional helper functions that allow you to convert between wide strings and narrow strings using the provided string converters:

C++

```
enum ConversionFlags { strictConversion, lenientConversion };

std::string wstringToString(const std::wstring&, const StringConverterPtr& = 0,
                const WstringConverterPtr& = 0,
ConversionFlags = lenientConversion);
std::wstring stringToWstring(const std::string&, const StringConverterPtr& = 0,
                const WstringConverterPtr& = 0,
ConversionFlags = lenientConversion);
```

When the narrow string converter given to `wstringToString` is null, the encoding of the returned narrow string is UTF-8. When the wide string converter given to `wstringToString` is null, the encoding of wide string parameter is UTF-16 or UTF-32, depending on the size of `wchar_t`.

Likewise for `stringToWstring`, when the wide string converter is null, the encoding of the returned wide string is UTF-16 or UTF-32 depending on the size of wchar_t. When the narrow string converter given to `stringToWstring` is null, the encoding of narrow string parameter is UTF-8.

When using the built-in converters (through null parameters), byte sequences that are illegal, such as `0xF4908080`, result in a `UTFConversionException`. For other errors, the `ConversionFlags` parameter determines how rigorously the functions check for errors. When set to `lenientConversion` (the default), the functions tolerate isolated surrogates and irregular sequences, and substitute the UTF-32 replacement character `0x0000FFFD` for character values above `0x10FFFF`. When set to `strictConversion`, the functions do not tolerate such errors and throw a `UTFConversionException` instead:

Copyright 2016, ZeroC, Inc.

C++

```
enum ConversionError { partialCharacter, badEncoding };

class UTFConversionException : public Exception {
public:
    UTFConversionException(const char* file, int line, ConversionError r);

    ConversionError conversionError() const;
    // ...
};
```

The `conversionError` member function returns the reason for the failure:

- `partialCharacter`
  The UTF-8 source string contains a trailing incomplete UTF-8 byte sequence.
- `badEncoding`
  The UTF-8 source string contains a byte sequence that is not a valid UTF-8 encoded character, or the Unicode source string contains a bit pattern that does not represent a valid Unicode character.

See Also

- UTF-8 Conversion

Copyright 2016, ZeroC, Inc.

The iconv String Converter

For Linux and Unix platforms, `IceUtil` provides an `IconvStringConverter` template class that uses the `iconv` conversion facility to convert between the native encoding and UTF-8. The only member function of interest is the constructor:

C++

```
template<typename charT>
class IconvStringConverter : public IceUtil::BasicStringConverter<charT>
{
public:
    IconvStringConverter(const char* = nl_langinfo(CODESET));

    // ...
};
```

To use this string converter, you specify whether the conversion you want is for narrow or wide characters via the template argument, and you specify the corresponding native encoding with the constructor argument. For example, to create a converter that converts between ISO Latin-1 and UTF-8, you can instantiate the converter as follows:

C++

```
StringConverterPtr stringConverter = new IconvStringConverter<char>("ISO-8859-1");
```

Similarly, to convert between the internal wide character encoding and UTF-8, you can instantiate a converter as follows:

C++

```
WstringConverterPtr
wstringConverter = new IconvStringConverter<wchar_t>("WCHAR_T");
```

The string you pass to the constructor must be one of the values returned by `iconv -l`, which lists all the available character encodings for your machine.

Using the `IconvStringConverter` template makes it easy to install code converters for any available encoding without having to explicitly write (or call) conversion routines, whose implementation is typically non-trivial.

See Also

- Installing String Converters

Copyright 2016, ZeroC, Inc.

The Ice String Converter Plug-in

The Ice run time includes a plug-in that supports conversion between UTF-8 and native encodings on Unix and Windows platforms. You can use this plug-in to install converters for narrow and wide strings into the communicator of an existing program. This feature is primarily intended for use in scripting language extensions such as Ice for Python; if you need to use string converters in your C++ application, we recommend using the technique described in Installing String Converters instead.

Note that an application must be designed to operate correctly in the presence of a string converter. A string converter assumes that it converts strings in the native encoding into the UTF-8 encoding, and vice versa. An application that performs its own conversions on strings that cross a Slice interface boundary can cause encoding errors when those strings are processed by a converter.

Configuring the Ice String Converter Plug-in

You can install the plug-in using a configuration property like the one shown below:

```
Ice.Plugin.IceStringConverter=Ice:createStringConverter iconv=encoding[,encoding]
windows=code-page
```

The first component of the property value represents the plug-in's entry point, which includes the abbreviated name of the shared library or DLL (`Ice`) and the name of a factory function (`createStringConverter`).

The plug-in accepts the following arguments:

- `iconv=encoding[,encoding]`
  This argument is optional on Unix platforms and ignored on Windows platforms. If specified, it defines the `iconv` names of the narrow string encoding and the optional wide-string encoding. If this argument is not specified, the plug-in installs a narrow string converter that uses the default locale-dependent encoding.
- `windows=code-page`
  This argument is required on Windows platforms and ignored on Unix platforms. The *`code-page`* value represents a code page number, such as `1252`.

The plug-in's argument semantics are designed so that the same configuration property can be used on both Windows and Unix platforms, as shown in the following example:

```
Ice.Plugin.IceStringConverter=Ice:createStringConverter
iconv=ISO8859-1 windows=1252
```

If the configuration file containing this property is shared by programs in multiple implementation languages, you can use an alternate syntax that is loaded only by the Ice for C++ run time:

```
Ice.Plugin.IceStringConverter.cpp=Ice:createStringConverter iconv=ISO8859-1
windows=1252
```

If using Ice-E, you must call the `Ice::registerIceStringConverter()` function to ensure the plugin is linked with your application.

`Ice::registerIceStringConverter` is a simple helper function that calls `Ice::registerPluginFactory`.

See Also

- UTF-8 Conversion
- Installing String Converters
- Plug-in Configuration
- Ice.InitPlugins
- Ice.Plugin.*
- Ice.PluginLoadOrder

Copyright 2016, ZeroC, Inc.

Custom String Converter Plug-ins

If the default string converter plug-in does not satisfy your requirements, you can install your own string converters with a plugin, for example:

C++

```
class MyStringConverterPlugin : public Ice::Plugin {
public:

    MyStringConverterPlugin(const IceUtil::StringConverterPtr& stringConverter,
                            const IceUtil::WstringConverterPtr& wstringConverter
= 0)
    {
        IceUtil::setProcessStringConverter(stringConverter);
        IceUtil::setProcessWstringConverter(wstringConverter);
    }

    virtual void initialize() {}
    virtual void destroy() {}
};
}
```

Like in this example, you should install the string converters in your plugin's constructor. Do not install the string converters in `initialize`.

In order to create such a plug-in, you must do the following:

- Define and export a factory function that returns an instance of your plugin class
- Implement the string converter(s) that you will pass to your plugin's constructor, or use the ones included with Ice.
- Package your code into a shared library or DLL.

To install your plug-in, use a configuration property like the one shown below:

```
Ice.Plugin.MyConverterPlugin=myconverter:createConverter ...
```

The first component of the property value represents the plug-in's entry point, which includes the abbreviated name of the shared library or DLL (`myconverter`) and the name of a factory function (`createConverter`).

If the configuration file containing this property is shared by programs in multiple implementation languages, you can use an alternate syntax that is loaded only by the Ice for C++ run time:

```
Ice.Plugin.MyConverterPlugin.cpp=myconverter:createConverter ...
```

See Also

- The Ice String Converter Plug-in
- Installing String Converters
- Plug-in API
- Ice.Plugin.*
- Ice.InitPlugins
- Ice.PluginLoadOrder

Copyright 2016, ZeroC, Inc.

The C++ Utility Library

Ice for C++ includes a number of utility classes and functions in the `IceUtil` namespace, which we summarize here for your reference. Many of the classes and functions in `IceUtil` are documented elsewhere in this manual so, where appropriate, the sections here simply reference the relevant pages.

Topics

- Threads and Concurrency with C++
- The C++ AbstractMutex Class
- The C++ Cache Template
- The C++ Exception Class
- The C++ generateUUID Function
- The C++ Handle Template
- The C++ Handle Template Adaptors
- The C++ ScopedArray Template
- The C++ Shared and SimpleShared Classes
- The C++ Time Class
- The C++ Timer and TimerTask Classes

Copyright 2016, ZeroC, Inc.

Threads and Concurrency with C++

Threading and concurrency control vary widely with different operating systems. To make threads programming easier and portable, Ice provides a simple thread abstraction layer that allows you to write portable source code regardless of the underlying platform.

This section looks at the threading and concurrency control mechanisms in Ice for C++. It explains the threading abstractions provided by Ice: mutexes, monitors, and threads. Using these APIs allows you to make your code thread safe and to create threads of your own without having to use non-portable APIs that differ in syntax or semantics across different platforms: Ice not only provides a portable API but also guarantees that the semantics of the various functions are the same across different platforms. This makes it easier to create thread-safe applications and lets you move your code between platforms with simple recompilation.

This section assumes you are familiar with light-weight threads and concurrency control. Also see The Ice Threading Model, which provides a language-neutral introduction to the Ice threading model.

Library Overview

The Ice threading library provides the following thread-related abstractions:

- mutexes
- recursive mutexes
- monitors
- a thread abstraction that allows you to create, control, and destroy threads

The synchronization primitives permit you to implement concurrency control at different levels of granularity. In addition, the thread abstraction allows you to, for example, create a separate thread that can respond to GUI or other asynchronous events. All of the threading APIs are part of the `IceUtil` namespace.

Topics

- The C++ Mutex Class
- The C++ RecMutex Class
- The C++ Monitor Class
- The C++ Cond Class
- The C++ Thread Classes
- Priority Inversion in C++
- Portable Signal Handling in C++

Copyright 2016, ZeroC, Inc.

The C++ Mutex Class

This page describes how to use mutexes — one of the available synchronization primitives.

On this page:

- Mutex Member Functions
- Adding Thread Safety to the File System Application in C++
- Guaranteed Unlocking of Mutexes in C++

Mutex Member Functions

The class `IceUtil::Mutex` (defined in `IceUtil/Mutex.h`) provides a simple non-recursive mutual exclusion mechanism:

C++

```
namespace IceUtil {
    enum MutexProtocol { PrioInherit, PrioNone };

    class Mutex {
    public:
        Mutex();
        Mutex(MutexProtocol p);
        ~Mutex();

        void lock() const;
        bool tryLock() const;
        void unlock() const;

        typedef LockT<Mutex> Lock;
        typedef TryLockT<Mutex> TryLock;
    };
}
```

The member functions of this class work as follows:

- `Mutex()`
  `Mutex(MutexProtocol p)`
  You can optionally specify a mutex protocol when you construct a mutex. The mutex protocol controls how the mutex behaves with respect to thread priorities. Default-constructed mutexes use a system-wide default.
- `lock`
  The `lock` function attempts to acquire the mutex. If the mutex is already locked, it suspends the calling thread until the mutex becomes available. The call returns once the calling thread has acquired the mutex.
- `tryLock`
  The `tryLock` function attempts to acquire the mutex. If the mutex is available, the call returns true with the mutex locked. Otherwise, if the mutex is locked by another thread, the call returns false.
- `unlock`
  The `unlock` function unlocks the mutex.

Note that `IceUtil::Mutex` is a non-recursive mutex implementation. This means that you must adhere to the following rules:

- Do not call `lock` on the same mutex more than once from a thread. The mutex is not recursive so, if the owner of a mutex attempts to lock it a second time, the behavior is undefined.
- Do not call `unlock` on a mutex unless the calling thread holds the lock. Calling `unlock` on a mutex that is not currently held by any thread, or calling `unlock` on a mutex that is held by a different thread, results in undefined behavior.

Use the `IceUtil::RecMutex` class if you need recursive semantics.

Copyright 2016, ZeroC, Inc.

Adding Thread Safety to the File System Application in C++

Recall that the implementation of the `read` and `write` operations for our file system server is not thread safe:

C++

```
Filesystem::Lines
Filesystem::FileI::read(const Ice::Current&) const
{
    return _lines;      // Not thread safe!
}

void
Filesystem::FileI::write(const Filesystem::Lines& text,
                         const Ice::Current&)
{
    _lines = text;      // Not thread safe!
}
```

The problem here is that, if we receive concurrent invocations of `read` and `write`, one thread will be assigning to the `_lines` vector while another thread is reading that same vector. The outcome of such concurrent data access is undefined; to avoid the problem, we need to serialize access to the `_lines` member with a mutex. We can make the mutex a data member of the `FileI` class and lock and unlock it in the `read` and `write` operations:

Copyright 2016, ZeroC, Inc.

C++

```
#include <IceUtil/Mutex.h>
// ...

namespace Filesystem {
    // ...

    class FileI : virtual public File,
                  virtual public Filesystem::NodeI {
    public:
        // As before...
    private:
        Lines _lines;
        IceUtil::Mutex _fileMutex;
    };
    // ...
}

Filesystem::Lines
Filesystem::FileI::read(const Ice::Current&) const
{
    _fileMutex.lock();
    Lines l = _lines;
    _fileMutex.unlock();
    return l;
}

void
Filesystem::FileI::write(const Filesystem::Lines& text, const Ice::Current&)
{
    _fileMutex.lock();
    _lines = text;
    _fileMutex.unlock();
}
```

The `FileI` class here is identical to the original implementation, except that we have added the `_fileMutex` data member. The `read` and `write` operations lock and unlock the mutex to ensure that only one thread can read or write the file at a time. Note that, by using a separate mutex for each `FileI` instance, it is still possible for multiple threads to concurrently read or write files, as long as they each access a *different* file. Only concurrent accesses to the *same* file are serialized.

The implementation of `read` is somewhat awkward here: we must make a local copy of the file contents while we are holding the lock and return that copy. Doing so is necessary because we must unlock the mutex before we can return from the function. However, as we will see in the next section, the copy can be avoided by using a helper class that unlocks the mutex automatically when the function returns.

Guaranteed Unlocking of Mutexes in C++

Using the raw `lock` and `unlock` operations on mutexes has an inherent problem: if you forget to unlock a mutex, your program will deadlock. Forgetting to unlock a mutex is easier than you might suspect, for example:

Copyright 2016, ZeroC, Inc.

C++

```
Filesystem::Lines
Filesystem::File::read(const Ice::Current&) const
{
    _fileMutex.lock();                  // Lock the mutex
    Lines l = readFileContents();       // Read from database
    _fileMutex.unlock();                // Unlock the mutex
    return l;
}
```

Assume that we are keeping the contents of the file on secondary storage, such as a database, and that the `readFileContents` function accesses the file. The code is almost identical to the previous example but now contains a latent bug: if `readFileContents` throws an exception, the `read` function terminates without ever unlocking the mutex. In other words, this implementation of `read` is not exception-safe.

The same problem can easily arise if you have a larger function with multiple return paths. For example:

C++

```
void
SomeClass::someFunction(/* params here... */)
{
    _mutex.lock();                      // Lock a mutex

    // Lots of complex code here...

    if (someCondition) {
        // More complex code here...
        return;                         // Oops!!!
    }

    // More code here...

    _mutex.unlock();                    // Unlock the mutex
}
```

In this example, the early return from the middle of the function leaves the mutex locked. Even though this example makes the problem quite obvious, in large and complex pieces of code, both exceptions and early returns can cause hard-to-track deadlock problems. To avoid this, the `Mutex` class contains two type definitions for helper classes, called `Lock` and `TryLock`:

Copyright 2016, ZeroC, Inc.

C++

```
namespace IceUtil {

    class Mutex {
        // ...

        typedef LockT<Mutex> Lock;
        typedef TryLockT<Mutex> TryLock;
    };
}
```

`LockT` and `TryLockT` are simple templates that primarily consist of a constructor and a destructor; the `LockT` constructor calls `lock` on its argument, and the `TryLockT` constructor calls `tryLock` on its argument. The destructors call `unlock` if the mutex is locked when the template goes out of scope. By instantiating a local variable of type `Lock` or `TryLock`, we can avoid the deadlock problem entirely:

C++

```
void
SomeClass::someFunction(/* params here... */)
{
    IceUtil::Mutex::Lock lock(_mutex);  // Lock a mutex

    // Lots of complex code here...

    if (someCondition) {
        // More complex code here...
        return;                         // No problem
    }

    // More code here...

}   // Destructor of lock unlocks the mutex
```

This is an example of the *RAII* (*Resource Acquisition Is Initialization*) idiom [1].

On entry to `someFunction`, we instantiate a local variable `lock`, of type `IceUtil::Mutex::Lock`. The constructor of `lock` calls `lock` on the mutex so the remainder of the function is inside a critical region. Eventually, `someFunction` returns, either via an ordinary return (in the middle of the function or at the end) or because an exception was thrown somewhere in the function body. Regardless of how the function terminates, the C++ run time unwinds the stack and calls the destructor of `lock`, which unlocks the mutex, so we cannot get trapped by the deadlock problem we had previously.

Both the `Lock` and `TryLock` templates have a few member functions:

- `void acquire() const`
  This function attempts to acquire the lock and blocks the calling thread until the lock becomes available. If the caller calls `acquire` on a mutex it has locked previously, the function throws `ThreadLockedException`.
- `bool tryAcquire() const`
  This function attempts to acquire the mutex. If the mutex can be acquired, it returns true with the mutex locked; if the mutex cannot be acquired, it returns false. If the caller calls `tryAcquire` on a mutex it has locked previously, the function throws `ThreadLocked`

Copyright 2016, ZeroC, Inc.

`Exception`.

- `void release() const`
  This function releases a previously locked mutex. If the caller calls release on a mutex it has unlocked previously, the function throws `ThreadLockedException`.
- `bool acquired() const`
  This function returns true if the caller has locked the mutex previously, otherwise it returns false. If you use the `TryLock` template, you must call `acquired` after instantiating the template to test whether the lock actually was acquired.

These functions are useful if you want to use the `Lock` and `TryLock` templates for guaranteed unlocking, but need to temporarily release the lock:

C++

```
{
    IceUtil::Mutex::TryLock m(someMutex);

    if (m.acquired())
    {

        // Got the lock, do processing here...

        if (release_condition) {
            m.release();
        }

        // Mutex is now unlocked, someone else can lock it.
        // ...

        m.acquire(); // Block until mutex becomes available.

        // ...

        if (release_condition) {
            m.release();
        }

        // Mutex is now unlocked, someone else can lock it.

        // ...

        // Spin on the mutex until it becomes available.
        while (!m.tryLock()) {
            // Do some other processing here...
        }

        // Mutex locked again at this point.

        // ...
    }

} // Close scope, m is unlocked by its destructor.
```

Tip

Copyright 2016, ZeroC, Inc.

You should make it a habit to always use the `Lock` and `TryLock` helpers instead of calling `lock` and `unlock` directly. Doing so results in code that is easier to understand and maintain.

Using the `Lock` helper, we can rewrite the implementation of our `read` and `write` operations as follows:

C++

```
Filesystem::Lines
Filesystem::FileI::read(const Ice::Current&) const
{
    IceUtil::Mutex::Lock lock(_fileMutex);
    return _lines;
}

void
Filesystem::FileI::write(const Filesystem::Lines& text, const Ice::Current&)
{
    IceUtil::Mutex::Lock lock(_fileMutex);
    _lines = text;
}
```

Note that this also eliminates the need to make a copy of the `_lines` data member: the return value is initialized under protection of the mutex and cannot be modified by another thread once the destructor of `lock` unlocks the mutex.

See Also

- Example of a File System Server in C++
- Priority Inversion in C++
- The C++ RecMutex Class

References

1. Stroustrup, B. 1997. *The C++ Programming Language*. Reading, MA: Addison-Wesley.

Copyright 2016, ZeroC, Inc.

The C++ RecMutex Class

A non-recursive mutex cannot be locked more than once, even by the thread that holds the lock. This frequently becomes a problem if a program contains a number of functions, each of which must acquire a mutex, and you want to call one function as part of the implementation of another function:

C++

```
IceUtil::Mutex _mutex;

void
f1()
{
    IceUtil::Mutex::Lock lock(_mutex);
    // ...
}

void
f2()
{
    IceUtil::Mutex::Lock lock(_mutex);
    // Some code here...

    // Call f1 as a helper function
    f1();                               // Deadlock!

    // More code here...
}
```

`f1` and `f2` each correctly lock the mutex before manipulating data but, as part of its implementation, `f2` calls `f1`. At that point, the program deadlocks because `f2` already holds the lock that `f1` is trying to acquire. For this simple example, the problem is obvious. However, in complex systems with many functions that acquire and release locks, it can get very difficult to track down this kind of situation: the locking conventions are not manifest anywhere but in the source code and each caller must know which locks to acquire (or not to acquire) before calling a function. The resulting complexity can quickly get out of hand.

Ice provides a recursive mutex class `RecMutex` (defined in `IceUtil/RecMutex.h`) that avoids this problem:

Copyright 2016, ZeroC, Inc.

C++

```
namespace IceUtil {

    class RecMutex {
    public:
        RecMutex();
        RecMutex(MutexProtocol p);
        ~RecMutex();

        void lock() const;
        bool tryLock() const;
        void unlock() const;

        typedef LockT<RecMutex> Lock;
        typedef TryLockT<RecMutex> TryLock;
    };
}
```

Note that the signatures of the operations are the same as for `IceUtil::Mutex`. However, `RecMutex` implements a recursive mutex:

- `RecMutex()`
  `RecMutex(MutexProtocol p)`
  You can optionally specify a mutex protocol when you construct a mutex. The mutex protocol controls how the mutex behaves with respect to thread priorities. Default-constructed mutexes use a system-wide default.
- `lock`
  The `lock` function attempts to acquire the mutex. If the mutex is already locked by another thread, it suspends the calling thread until the mutex becomes available. If the mutex is available or is already locked by the calling thread, the call returns immediately with the mutex locked.
- `tryLock`
  The `tryLock` function works like `lock`, but, instead of blocking the caller, it returns false if the mutex is locked by another thread. Otherwise, the return value is true.
- `unlock`
  The `unlock` function unlocks the mutex.

As for non-recursive mutexes, you must adhere to a few simple rules for recursive mutexes:

- Do not call `unlock` on a mutex unless the calling thread holds the lock.
- You must call `unlock` as many times as you called `lock` for the mutex to become available to another thread. (Internally, a recursive mutex is implemented with a counter that is initialized to zero. Each call to `lock` increments the counter and each call to `unlock` decrements the counter; the mutex is made available to another thread when the counter returns to zero.)

Using recursive mutexes, the code fragment shown earlier works correctly:

Copyright 2016, ZeroC, Inc.

C++

```
#include <IceUtil/RecMutex.h>
// ...

IceUtil::RecMutex _mutex;       // Recursive mutex

void
f1()
{
    IceUtil::RecMutex::Lock lock(_mutex);
    // ...
}

void
f2()
{
    IceUtil::RecMutex::Lock lock(_mutex);
    // Some code here...

    // Call f1 as a helper function
    f1();                               // Fine

    // More code here...
}
```

Note that the type of the mutex is now `RecMutex` instead of `Mutex`, and that we are using the `Lock` type definition provided by the `RecMutex` class, not the one provided by the `Mutex` class.

See Also

- The C++ Mutex Class
- Priority Inversion in C++

Copyright 2016, ZeroC, Inc.

The C++ Monitor Class

The recursive and non-recursive mutex classes implement a simple mutual exclusion mechanism that allows only a single thread to be active in a critical region at a time. In particular, for a thread to enter the critical region, another thread must leave it. This means that, with mutexes, it is impossible to suspend a thread inside a critical region and have that thread wake up again at a later time, for example, when a condition becomes true.

To address this problem, Ice provides a monitor. Briefly, a monitor is a synchronization mechanism that protects a critical region: as for a mutex, only one thread may be active at a time inside the critical region. However, a monitor allows you to suspend a thread inside the critical region; doing so allows another thread to enter the critical region. The second thread can either leave the monitor (thereby unlocking the monitor), or it can suspend itself inside the monitor; either way, the original thread is woken up and continues execution inside the monitor. This extends to any number of threads, so several threads can be suspended inside a monitor.

The monitors provided by Ice have *Mesa* semantics, so called because they were first implemented by the Mesa programming language [1]. Mesa monitors are provided by a number of languages, including Java and Ada. With Mesa semantics, the signalling thread continues to run and another thread gets to run only once the signalling thread suspends itself or leaves the monitor.

Monitors provide a more flexible mutual exclusion mechanism than mutexes because they allow a thread to check a condition and, if the condition is false, put itself to sleep; the thread is woken up by some other thread that has changed the condition.

On this page:

- Monitor Member Functions
- Using Monitors in C++
- Efficient Notification using Monitors in C++

Monitor Member Functions

Ice provides monitors with the `IceUtil::Monitor` class (defined in `IceUtil/Monitor.h`):

C++

```
namespace IceUtil {

    template <class T>
    class Monitor {
    public:
        void lock() const;
        void unlock() const;
        bool tryLock() const;

        void wait() const;
        bool timedWait(const Time&) const;
        void notify();
        void notifyAll();

        typedef LockT<Monitor<T> > Lock;
        typedef TryLockT<Monitor<T> > TryLock;
    };
}
```

Note that `Monitor` is a template class that requires either `Mutex` or `RecMutex` as its template parameter. (Instantiating a `Monitor` with a `RecMutex` makes the monitor recursive.)

The member functions behave as follows:

- `lock`
  This function attempts to lock the monitor. If the monitor is currently locked by another thread, the calling thread is suspended until the monitor becomes available. The call returns with the monitor locked.

Copyright 2016, ZeroC, Inc.

- `tryLock`
  This function attempts to lock a monitor. If the monitor is available, the call returns true with the monitor locked. If the monitor is locked by another thread, the call returns false.
- `unlock`
  This function unlocks a monitor. If other threads are waiting to enter the monitor (are blocked inside a call to `lock`), one of the threads is woken up and locks the monitor.
- `wait`
  This function suspends the calling thread and, at the same time, releases the lock on the monitor. A thread suspended inside a call to `wait` can be woken up by another thread that calls `notify` or `notifyAll`. When the call returns, the suspended thread resumes execution with the monitor locked.
- `timedWait`
  This function suspends the calling thread for up to the specified timeout. If another thread calls `notify` or `notifyAll` and wakes up the suspended thread before the timeout expires, the call returns true and the suspended thread resumes execution with the monitor locked. Otherwise, if the timeout expires, the function returns false. Wait intervals are represented by instances of the `Time` class.
- `notify`
  This function wakes up a single thread that is currently suspended in a call to `wait` or `timedWait`. If no thread is suspended in a call to `wait` or `timedWait` at the time `notify` is called, the notification is lost (that is, calls to `notify` are *not* remembered if there is no thread to be woken up). Note that notifying does not run another thread immediately. Another thread gets to run only once the notifying thread either calls `wait` or `timedWait` or unlocks the monitor (Mesa semantics).
- `notifyAll`
  This function wakes up all threads that are currently suspended in a call to `wait` or `timedWait`. As for `notify`, calls to `notifyAll` are lost if no threads are suspended at the time. Also as for `notify`, `notifyAll` causes other threads to run only once the notifying thread has either called `wait` or `timedWait` or unlocked the monitor (Mesa semantics).

You must adhere to a few rules for monitors to work correctly:

- Do not call `unlock` unless you hold the lock. If you instantiate a monitor with a recursive mutex, you get recursive semantics, that is, you must call `unlock` as many times as you have called `lock` (or `tryLock`) for the monitor to become available.
- Do not call `wait` or `timedWait` unless you hold the lock.
- Do not call `notify` or `notifyAll` unless you hold the lock.
- When returning from a `wait` call, you *must* re-test the condition before proceeding (as shown below).

Using Monitors in C++

To illustrate how to use a monitor, consider a simple unbounded queue of items. A number of producer threads add items to the queue, and a number of consumer threads remove items from the queue. If the queue becomes empty, consumers must wait until a producer puts a new item on the queue. The queue itself is a critical region, that is, we cannot allow a producer to put an item on the queue while a consumer is removing an item. Here is a very simple implementation of a such a queue:

Copyright 2016, ZeroC, Inc.

C++

```
template<class T> class Queue {
public:
    void put(const T& item) {
        _q.push_back(item);
    }

    T get() {
        T item = _q.front();
        _q.pop_front();
        return item;
    }

private:
    list<T> _q;
};
```

As you can see, producers call the `put` method to enqueue an item, and consumers call the `get` method to dequeue an item. Obviously, this implementation of the queue is not thread-safe and there is nothing to stop a consumer from attempting to dequeue an item from an empty queue.

Here is a version of the queue that uses a monitor to suspend a consumer if the queue is empty:

C++

```
#include <IceUtil/Monitor.h>

template<class T> class Queue : public IceUtil::Monitor<IceUtil::Mutex> {
public:
    void put(const T& item) {
        IceUtil::Monitor<IceUtil::Mutex>::Lock lock(*this);
        _q.push_back(item);
        notify();
    }

    T get() {
        IceUtil::Monitor<IceUtil::Mutex>::Lock lock(*this);
        while (_q.size() == 0)
            wait();
        T item = _q.front();
        _q.pop_front();
        return item;
    }

private:
    list<T> _q;
};
```

Note that the `Queue` class now inherits from `IceUtil::Monitor<IceUtil::Mutex>`, that is, `Queue` *is-a* monitor.

Copyright 2016, ZeroC, Inc.

Both the `put` and `get` methods lock the monitor when they are called. As for mutexes, instead of calling `lock` and `unlock` directly, we are using the `Lock` helper which automatically locks the monitor when it is instantiated and unlocks the monitor again when it is destroyed.

The `put` method first locks the monitor and then, now being in sole possession of the critical region, enqueues an item. Before returning (thereby unlocking the monitor), `put` calls `notify`. The call to `notify` will wake up any consumer thread that may be asleep in a `wait` call to inform the consumer that an item is available.

The `get` method also locks the monitor and then, before attempting to dequeue an item, tests whether the queue is empty. If so, the consumer calls `wait`. This suspends the consumer inside the `wait` call and unlocks the monitor, so a producer can enter the monitor to enqueue an item. Once that happens, the producer calls `notify`, which causes the consumer's `wait` call to complete, with the monitor again locked for the consumer. The consumer now dequeues an item and returns (thereby unlocking the monitor).

For this machinery to work correctly, the implementation of `get` does two things:

- `get` tests whether the queue is empty *after* acquiring the lock.
- `get` re-tests the condition in a loop around the call to `wait`; if the queue is still empty after `wait` returns, the `wait` call is re-entered.

You *must* always write your code to follow the same pattern:

- *Never* test a condition unless you hold the lock.
- *Always* re-test the condition in a loop around `wait`. If the test still shows the wrong outcome, call `wait` again.

Not adhering to these conditions will eventually result in a thread accessing shared data when it is not in its expected state, for the following reasons:

1. If you test a condition without holding the lock, there is nothing to prevent another thread from entering the monitor and changing its state before you can acquire the lock. This means that, by the time you get around to locking the monitor, the state of the monitor may no longer be in agreement with the result of the test.
2. Some thread implementations suffer from a problem known as *spurious wake-up*: occasionally, more than one thread may wake up in response to a call to `notify`, or a thread may wake up without any call to `notify` at all. As a result, each thread that returns from a call to `wait` must re-test the condition to ensure that the monitor is in its expected state: the fact that `wait` returns does *not* i ndicate that the condition has changed.

Efficient Notification using Monitors in C++

The previous implementation of our thread-safe queue unconditionally notifies a waiting reader whenever a writer deposits an item into the queue. If no reader is waiting, the notification is lost and does no harm. However, unless there is only a single reader and writer, many notifications will be sent unnecessarily, causing unwanted overhead.

Here is one way to fix the problem:

Copyright 2016, ZeroC, Inc.

C++

```
#include <IceUtil/Monitor.h>

template<class T> class Queue : public IceUtil::Monitor<IceUtil::Mutex> {
public:
    void put(const T& item) {
        IceUtil::Monitor<IceUtil::Mutex>::Lock lock(*this);
        _q.push_back(item);
        if (_q.size() == 1)
            notify();
    }

    T get() {
        IceUtil::Monitor<IceUtil::Mutex>::Lock lock(*this);
        while (_q.size() == 0)
            wait();
        T item = _q.front();
        _q.pop_front();
        return item;
    }

private:
    list<T> _q;
};
```

The only difference between this code and the implementation shown earlier is that a writer calls `notify` only if the queue length has just changed from empty to non-empty. That way, unnecessary `notify` calls are never made. However, this approach works only for a single reader thread. To see why, consider the following scenario:

1. Assume that the queue currently contains a number of items and that we have five reader threads.
2. The five reader threads continue to call `get` until the queue becomes empty and all five readers are waiting in `get`.
3. The scheduler schedules a writer thread. The writer finds the queue empty, deposits an item, and wakes up a single reader thread.
4. The awakened reader thread dequeues the single item on the queue.
5. The reader calls `get` a second time, finds the queue empty, and goes to sleep again.

The net effect of this is that there is a good chance that only one reader thread will ever be active; the other four reader threads end up being permanently asleep inside the `get` method.

One way around this problem is call `notifyAll` instead of `notify` once the queue length exceeds a certain amount, for example:

Copyright 2016, ZeroC, Inc.

C++

```
#include <IceUtil/Monitor.h>

template<class T> class Queue : public IceUtil::Monitor<IceUtil::Mutex> {
public:
    void put(const T& item) {
        IceUtil::Monitor<IceUtil::Mutex>::Lock lock(*this);
        _q.push_back(item);
        if (_q.size() >= _wakeupThreshold)
            notifyAll();
    }

    T get() {
        IceUtil::Monitor<IceUtil::Mutex>::Lock lock(*this);
        while (_q.size() == 0)
            wait();
        T item = _q.front();
        _q.pop_front();
        return item;
    }

private:
    list<T> _q;
    const int _wakeupThreshold = 100;
};
```

Here, we have added a private data member `_wakeupThreshold`; a writer wakes up *all* waiting readers once the queue length exceeds the threshold, in the expectation that all the readers will consume items more quickly than they are produced, thereby reducing the queue length below the threshold again.

This approach works, but has drawbacks as well:

- The appropriate value of `_wakeupThreshold` is difficult to determine and sensitive to things such as speed and number of processors and I/O bandwidth.
- If multiple readers are asleep, they are all made runnable by the thread scheduler once a writer calls `notifyAll`. On a multiprocessor machine, this may result in all readers running at once (one per CPU). However, as soon as the readers are made runnable, each of them attempts to reacquire the mutex that protects the monitor before returning from `wait`. Of course, only one of the readers actually succeeds and the remaining readers are suspended again, waiting for the mutex to become available. The net result is a large number of thread context switches as well as repeated and unnecessary locking of the system bus.

A better option than calling `notifyAll` is to wake up waiting readers one at a time. To do this, we keep track of the number of waiting readers and call `notify` only if a reader needs to be woken up:

Copyright 2016, ZeroC, Inc.

C++

```
#include <IceUtil/Monitor.h>

template<class T> class Queue : public IceUtil::Monitor<IceUtil::Mutex> {
public:
    Queue() : _waitingReaders(0) {}

    void put(const T& item) {
        IceUtil::Monitor<IceUtil::Mutex>::Lock lock(*this);
        _q.push_back(item);
        if (_waitingReaders)
            notify();
    }

    T get() {
        IceUtil::Monitor<IceUtil::Mutex>::Lock lock(*this);
        while (_q.size() == 0) {
            try {
                ++_waitingReaders;
                wait();
                --_waitingReaders;
            } catch (...) {
                --_waitingReaders;
                throw;
            }
        }
        T item = _q.front();
        _q.pop_front();
        return item;
    }

private:
    list<T> _q;
    short _waitingReaders;
};
```

This implementation uses a member variable `_waitingReaders` to keep track of the number of readers that are suspended. The constructor initializes the variable to zero and the implementation of `get` increments and decrements the variable around the call to `wait`. Note that these statements are enclosed in a `try-catch` block; this ensures that the count of waiting readers remains accurate even if `wait` throws an exception. Finally, `put` calls `notify` only if there is a waiting reader.

The advantage of this implementation is that it minimizes contention on the monitor mutex: a writer wakes up only a single reader at a time, so we do not end up with multiple readers simultaneously trying to lock the mutex. Moreover, the monitor `notify` implementation signals a waiting thread only *after* it has unlocked the mutex. This means that, when a thread wakes up from its call to `wait` and tries to reacquire the mutex, the mutex is likely to be unlocked. This results in more efficient operation because acquiring an unlocked mutex is typically very efficient, whereas forcefully putting a thread to sleep on a locked mutex is expensive (because it forces a thread context switch).

See Also

- The C++ Mutex Class
- The C++ RecMutex Class
- The C++ Time Class

References

1. Mitchell, J. G., et al. 1979. *Mesa Language Manual*. CSL-793. Palo Alto, CA: Xerox PARC.

Copyright 2016, ZeroC, Inc.

The C++ Cond Class

Condition variables are similar to monitors in that they allow a thread to enter a critical region, test a condition, and sleep inside the critical region while releasing its lock. Another thread then is free to enter the critical region, change the condition, and eventually signal the sleeping thread, which resumes at the point where it went to sleep and with the critical region once again locked.

Note that condition variables provide a subset of the functionality of monitors, so a monitor can always be used instead of a condition variable. However, condition variables are smaller, which may be important if you are seriously constrained with respect to memory.

Condition variables are provided by the `IceUtil::Cond` class. Here is its interface:

C++

```
class Cond : private noncopyable {
public:

    Cond();
    ~Cond();

    void signal();
    void broadcast();

    template<typename Lock>
    void wait(const Lock& lock) const;

    template<typename Lock>
    bool timedWait(const Lock& lock, const Time& timeout) const;
};
```

Using a condition variable is very similar to using a monitor. The main difference in the `Cond` interface is that the `wait` and `timedWait` me mber functions are template functions, instead of the entire class being a template. The member functions behave as follows:

- `wait`
  This function suspends the calling thread and, at the same time, releases the lock of the condition variable. A thread suspended inside a call to `wait` can be woken up by another thread that calls `signal` or `broadcast`. When `wait` completes, the suspended thread resumes execution with the lock held.
- `timedWait`
  This function suspends the calling thread for up to the specified timeout. If another thread calls `signal` or `broadcast` and wakes up the suspended thread before the timeout expires, the call returns true and the suspended thread resumes execution with the lock held. Otherwise, if the timeout expires, the function returns false. Wait intervals are represented by instances of the `Time` class.
- `signal`
  This function wakes up a single thread that is currently suspended in a call to `wait` or `timedWait`. If no thread is suspended in a call to `wait` or `timedWait` at the time `signal` is called, the signal is lost (that is, calls to `signal` are *not* remembered if there is no thread to be woken up). Note that signalling does not necessarily run another thread immediately; the thread calling `signal` may continue to run. However, depending on the underlying thread library, `signal` may also cause an immediate context switch to another thread.
- `broadcast`
  This function wakes up all threads that are currently suspended in a call to `wait` or `timedWait`. As for `signal`, calls to `broadcas t` are lost if no threads are suspended at the time.

You must adhere to a few rules for condition variables to work correctly:

- Do not call `wait` or `timedWait` unless you hold the lock.
- When returning from a `wait` call, you *must* re-test the condition before proceeding, just as for a monitor.

In contrast to monitors, which require you to call `notify` and `notifyAll` with the lock held, condition variables permit you to call `signal` a nd `broadcast` without holding the lock. Here is a code example that changes a condition and signals on a condition variable:

Copyright 2016, ZeroC, Inc.

C++

```
Mutex m;
Cond c;

// ...

{
     Mutex::Lock sync(m);

     // Change some condition other threads may be sleeping on...

     c.signal();

     // ...
} // m is unlocked here
```

This code is correct and will work as intended, but it is potentially inefficient. Consider the code executed by the waiting thread:

C++

```
{
    Mutex::Lock sync(m);

    while(!condition) {
        c.wait(sync);
    }

    // Condition is now true, do some processing...

} // m is unlocked here
```

Again, this code is correct and will work as intended. However, consider what can happen once the first thread calls `signal`. It is possible that the call to `signal` will cause an immediate context switch to the waiting thread. But, even if the thread implementation does not cause such an immediate context switch, it is possible for the signalling thread to be suspended after it has called `signal`, but before it unlocks the mutex `m`. If this happens, the following sequence of events occurs:

1. The waiting thread is still suspended inside the implementation of `wait` and is now woken up by the call to `signal`.
2. The now-awake thread tries to acquire the mutex `m` but, because the signalling thread has not yet released the mutex, is suspended again waiting for the mutex to be unlocked.
3. The signalling thread is scheduled again and leaves the scope enclosing `sync`, which unlocks the mutex, making the thread waiting for the mutex runnable.
4. The thread waiting for the mutex acquires the mutex and retests its condition.

While the preceding scenario is functionally correct, it is inefficient because it incurs two extra context switches between the signalling thread and the waiting thread. Because context switches are expensive, this can have quite a large impact on run-time performance, especially if the critical region is small and the condition changes frequently.

You can avoid the inefficiency by unlocking the mutex *before* calling signal:

Copyright 2016, ZeroC, Inc.

C++

```
Mutex m;
Cond c;

// ...

{
     Mutex::Lock sync(m);

     // Change some condition other threads may be sleeping on...

} // m is unlocked here

c.signal(); // Signal with the lock available
```

By arranging the code as shown, you avoid the additional context switches because, when the waiting thread is woken up by the call to `signal`, it succeeds in acquiring the mutex before returning from `wait` without being suspended and woken up again first.

As for monitors, you should exercise caution in using `broadcast`, particularly if you have many threads waiting on a condition. Condition variables suffer from the same potential problem as monitors with respect to `broadcast`, namely, that all threads that are currently suspended inside `wait` can immediately attempt to acquire the mutex, but only one of them can succeed and all other threads are suspended again. If your application is sensitive to this condition, you may want to consider waking threads in a more controlled manner.
See Also

- The C++ Monitor Class
- The C++ Time Class

Copyright 2016, ZeroC, Inc.

The C++ Thread Classes

The server-side Ice run time by default creates a thread pool for you and automatically dispatches each incoming request in its own thread. As a result, you usually only need to worry about synchronization among threads to protect critical regions when you implement a server. However, you may wish to create threads of your own. For example, you might need a dedicated thread that responds to input from a user interface. And, if you have complex and long-running operations that can exploit parallelism, you might wish to use multiple threads for the implementation of that operation.

Ice provides a simple thread abstraction that permits you to write portable source code regardless of the native threading platform. This shields you from the native underlying thread APIs and guarantees uniform semantics regardless of your deployment platform.

On this page:

- The C++ Thread Class
- Implementing Threads in C++
- Creating Threads in C++
- The C++ ThreadControl Class
- C++ Thread Example

The C++ `Thread` Class

The basic thread abstraction in Ice is provided by two classes, `ThreadControl` and `Thread` (defined in `IceUtil/Thread.h`):

Copyright 2016, ZeroC, Inc.

C++

```
namespace IceUtil {

    class Time;

    class ThreadControl {
    public:
#ifdef _WIN32
        typedef DWORD ID;
#else
        typedef pthread_t ID;
#endif

        ThreadControl();
#ifdef _WIN32
        ThreadControl(HANDLE, DWORD);
#else
        ThreadControl(explicit pthread_t);
#endif
        ID id() const;

        void join();
        void detach();

        static void sleep(const Time&);
        static void yield();

        bool operator==(const ThreadControl&) const;
        bool operator!=(const ThreadControl&) const;

    };

    class Thread :virtual public Shared {
    public:
        virtual void run() = 0;

        ThreadControl start(size_t stBytes = 0);
        ThreadControl start(size_t stBytes, int priority);
        ThreadControl getThreadControl() const;
        bool isAlive() const;

        bool operator==(const Thread&) const;
        bool operator!=(const Thread&) const;
        bool operator<(const Thread&) const;
    };
    typedef Handle<Thread> ThreadPtr;
}
```

The `Thread` class is an abstract base class with a pure virtual `run` method. To create a thread, you must specialize the `Thread` class and implement the `run` method (which becomes the starting stack frame for the new thread). Note that you must not allow any exceptions to escape from `run`. The Ice run time installs an exception handler that calls `::std::terminate` if `run` terminates with an exception.

Copyright 2016, ZeroC, Inc.

The remaining member functions behave as follows:

- `start(size_t stBytes = 0)`
  `start(size_t stBytes, int priority)`
  This member function starts a newly-created thread (that is, calls the `run` method). The `stBytes` parameter specifies a stack size (in bytes) for the thread. The default value of zero creates the thread with a default stack size that is determined by the operating system.

  You can also specify a priority for the thread. (If you do not supply a priority, the thread is created with the system default priority.) The priority value is system-dependent; on POSIX systems, the value must be a legal value for the `SCHED_RR` real-time scheduling policy. (`SCHED_RR` requires root privileges.) On Windows systems, the priority value is passed through to the Windows `setThreadPriority` function. Priority Inversion in C++ provides information about how you can deal with priority inversion.

  The return value is a `ThreadControl` object for the new thread.

  You can start a thread only once; calling `start` on an already-started thread raises `ThreadStartedException`.

  If the calling thread passes an invalid priority or, on POSIX systems, does not have root privileges, `start` raises `ThreadSyscallException`.

- `getThreadControl`
  This member function returns a `ThreadControl` object for the thread on which it is invoked. Calling this method before calling `start` raises a `ThreadNotStartedException`.

- `id`
  This method returns the underlying thread ID (`DWORD` for Windows and `pthread_t` for POSIX threads). This method is provided mainly for debugging purposes. Note also that `pthread_t` is, strictly-speaking, an opaque type, so you should not make any assumptions about what you can do with a thread ID.

- `isAlive`
  This method returns false before a thread's `start` method has been called and after a thread's `run` method has completed; otherwise, while the thread is still running, it returns true. `isAlive` is useful to implement a non-blocking join:

  C++

  ```
  ThreadPtr p = new MyThread();
  // ...
  while(p->isAlive()) {
      // Do something else...
  }
  p->getThreadControl().join(); // Will not block
  ```

- `operator==`
  `operator!=`
  `operator<`
  These member functions compare the in-memory address of two threads. They are provided so you can use `Thread` objects in sorted STL containers.

Note that `IceUtil` also defines the type `ThreadPtr`. This is the usual reference-counted smart pointer to guarantee automatic clean-up: the `Thread` destructor calls `delete this` once its reference count drops to zero.

Implementing Threads in C++

To illustrate how to implement threads, consider the following code fragment:

Copyright 2016, ZeroC, Inc.

C++

```
#include <IceUtil/Thread.h>
// ...

Queue q;

class ReaderThread : public IceUtil::Thread {
    virtual void run() {
        for (int i = 0; i < 100; ++i)
            cout << q.get() << endl;
    }
};

class WriterThread : public IceUtil::Thread {
    virtual void run() {
        for (int i = 0; i < 100; ++i)
            q.put(i);
    }
};
```

This code fragment defines two classes, `ReaderThread` and `WriterThread`, that inherit from `IceUtil::Thread`. Each class implements the pure virtual `run` method it inherits from its base class. For this simple example, a writer thread places the numbers from 1 to 100 into an instance of the thread-safe `Queue` class we defined in our discussion of monitors, and a reader thread retrieves 100 numbers from the queue and prints them to `stdout`.

Creating Threads in C++

To create a new thread, we simply instantiate the thread and call its `start` method:

C++

```
IceUtil::ThreadPtr t = new ReaderThread;
t->start();
// ...
```

Note that we assign the return value from `new` to a smart pointer of type `ThreadPtr`. This ensures that we do not suffer a memory leak:

1. When the thread is created, its reference count is set to zero.
2. Prior to calling `run` (which is called by the `start` method), `start` increments the reference count of the thread to 1.
3. For each `ThreadPtr` for the thread, the reference count of the thread is incremented by 1, and for each `ThreadPtr` that is destroyed, the reference count is decremented by 1.
4. When `run` completes, `start` decrements the reference count again and then checks its value: if the value is zero at this point, the `Thread` object deallocates itself by calling `delete this`; if the value is non-zero at this point, there are other smart pointers that reference this `Thread` object and deletion happens when the last smart pointer goes out of scope.

Note that, for all this to work, you *must* allocate your `Thread` objects on the heap — stack-allocated `Thread` objects will result in deallocation errors:

Copyright 2016, ZeroC, Inc.

C++

```
ReaderThread thread;
IceUtil::ThreadPtr t = &thread; // Bad news!!!
```

This is wrong because the destructor of `t` will eventually call `delete`, which has undefined behavior for a stack-allocated object.

Similarly, you *must* use a `ThreadPtr` for an allocated thread. Do not attempt to explicitly delete a thread:

C++

```
Thread* t = new ReaderThread();

// ...

delete t; // Disaster!
```

This will result in a double deallocation of the thread because the thread's destructor will call `delete this`.

It is legal for a thread to call `start` on itself from within its own constructor. However, if so, the thread must not be (very) short lived:

C++

```
class ActiveObject : public Thread() {
public:
    ActiveObject() {
        start();
    }

    void done() {
        getThreadControl().join();
    }

    virtual void run() {
        // *Very* short lived...
    }
};
typedef Handle<ActiveObject> ActiveObjectPtr;

// ...

ActiveObjectPtr ao = new ActiveObject;
```

With this code, it is possible for `run` to complete before the assignment to the smart pointer `ao` completes; in that case, `start` will call `delete this;` before it returns and `ao` ends up deleting an already-deleted object. However, note that this problem can arise only if `run` is indeed very short-lived and moreover, the scheduler allows the newly-created thread to run to completion before the assignment of the

Copyright 2016, ZeroC, Inc.

return value of `operator new` to `ao` takes place. This is highly unlikely to happen — if you are concerned about this scenario, do not call `start` from within a thread's own constructor. That way, the smart pointer is assigned first, and the thread started second, so the problem cannot arise.

The C++ `ThreadControl` Class

The `start` method returns an object of type `ThreadControl`. The member functions of `ThreadControl` behave as follows:

- `ThreadControl`
  The default constructor returns a `ThreadControl` object that refers to the calling thread. This allows you to get a handle to the current (calling) thread even if you had not previously saved a handle to that thread. For example:

  C++

  ```
  IceUtil::ThreadControl self;    // Get handle to self
  cout << self.id() << endl;      // Print thread ID
  ```

  This example also explains why we have two classes, `Thread` and `ThreadControl`: without a separate `ThreadControl`, it would not be possible to obtain a handle to an arbitrary thread. (Note that this code works even if the calling thread was not created by the Ice run time; for example, you can create a `ThreadControl` object for a thread that was created by the operating system.)

  The (implicit) copy constructor and assignment operator create a `ThreadControl` object that refers to the same underlying thread as the source `ThreadControl` object.

  Note that the constructor is overloaded. For Windows, the signature is

  C++

  ```
  ThreadControl(HANDLE, DWORD);
  ```

  For Unix, the signature is

  C++

  ```
  ThreadControl(pthread_t);
  ```

  These constructors allow you to create a `ThreadControl` object for the specified thread.

- `join`
  This method suspends the calling thread until the thread on which `join` is called has terminated. For example:

  C++

  ```
  IceUtil::ThreadPtr t = new ReaderThread; // Create a thread
  IceUtil::ThreadControl tc = t->start();  // Start it
  tc.join();                               // Wait for it
  ```


Copyright 2016, ZeroC, Inc.

If the reader thread has finished by the time the creating thread calls `join`, the call to `join` returns immediately; otherwise, the creating thread is suspended until the reader thread terminates.

Note that the `join` method of a thread must be called from only one other thread, that is, only one thread can wait for another thread to terminate. Calling `join` on a thread from more than one other thread has undefined behavior.

Calling `join` on a thread that was previously joined with or calling `join` on a detached thread has undefined behavior.
You must join with each thread you create; failure to join with a thread has undefined behavior.

- `detach`
  This method detaches a thread. Once a thread is detached, it cannot be joined with.

  Calling `detach` on an already detached thread, or calling `detach` on a thread that was previously joined with has undefined behavior.

  Note that, if you have detached a thread, you must ensure that the detached thread has terminated before your program leaves its `main` function. This means that, because detached threads cannot be joined with, they must have a life time that is shorter than that of the main thread.

- `sleep`
  This method suspends the calling thread for the amount of time specified by the `Time` class.

- `yield`
  This method causes the calling thread to relinquish the CPU, allowing another thread to run.

- `operator==`
  `operator!=`
  These operators compare thread IDs. (Note that `operator<` is not provided because it cannot be implemented portably.) These operators yield meaningful results only for threads that have not been detached or joined with.

As for all the synchronization primitives, you must adhere to a few rules when using threads to avoid undefined behavior:

- Do not allow `run` to throw an exception.
- Do not join with or detach a thread that you have not created yourself.
- For every thread you create, you must either join with that thread exactly once or detach it exactly once; failure to do so may cause resource leaks.
- Do not call `join` on a thread from more than one other thread.
- Do not leave `main` until all other threads you have created have terminated.
- Do not leave `main` until after you have destroyed all `Ice::Communicator` objects you have created (or use the `Ice::Application` class).
- A common mistake is to call `yield` from within a critical region. Doing so is usually pointless because the call to `yield` will look for another thread that can be run but, when that thread is run, it will most likely try to enter the critical region that is held by the yielding thread and go to sleep again. At best, this achieves nothing and, at worst, it causes many additional context switches for no gain. If you call `yield`, do so only in circumstances where there is at least a fair chance that another thread will actually be able to run and do something useful.

C++ Thread Example

Following is a small example that uses the `Queue` class we defined in our discussion of monitors. We create five writer and five reader threads. The writer threads each deposit 100 numbers into the queue, and the reader threads each retrieve 100 numbers and print them to `stdout`:

Copyright 2016, ZeroC, Inc.

C++

```
#include <vector>
#include <IceUtil/Thread.h>
// ...

Queue q;

class ReaderThread : public IceUtil::Thread {
    virtual void run() {
        for (int i = 0; i < 100; ++i)
            cout << q.get() << endl;
    }
};

class WriterThread : public IceUtil::Thread {
    virtual void run() {
        for (int i = 0; i < 100; ++i)
            q.put(i);
    }
};

int
main()
{
    vector<IceUtil::ThreadControl> threads;
    int i;

    // Create five reader threads and start them
    //
    for (i = 0; i < 5; ++i) {
        IceUtil::ThreadPtr t = new ReaderThread;
        threads.push_back(t->start());
    }

    // Create five writer threads and start them
    //
    for (i = 0; i < 5; ++i) {
        IceUtil::ThreadPtr t = new WriterThread;
        threads.push_back(t->start());
    }

    // Wait for all threads to finish
    //
    for (vector<IceUtil::ThreadControl>::iterator i = threads.begin();
         i != threads.end(); ++i) {
        i->join();
    }
}
```

The code uses the `threads` variable, of type `vector<IceUtil::ThreadControl>`, to keep track of the created threads. The code creates five reader and five writer threads, storing the `ThreadControl` object for each thread in the `threads` vector. Once all the threads

Copyright 2016, ZeroC, Inc.

are created and running, the code joins with each thread before returning from `main`.

Note that you must not leave `main` without first joining with the threads you have created: many threading libraries crash if you return from `main` with other threads still running. (This is also the reason why you must not terminate a program without first calling `Communicator::destroy`; the `destroy` implementation joins with all outstanding threads before it returns.)

See Also

- Smart Pointers for Classes
- The Server-Side main Function in C++
- The C++ Monitor Class
- The Ice Threading Model
- The C++ Time Class

Copyright 2016, ZeroC, Inc.

Priority Inversion in C++

In real-time systems, if you have threads with different priorities, it is possible to encounter a priority inversion. A priority inversion occurs when a low-priority thread prevents a higher-priority thread from running. This situation arises when a low-priority thread acquires a mutex and is pre-empted by one or more medium-priority threads that do not relinquish the CPU. If a high-priority thread then attempts to acquire the mutex locked by the low-priority thread, it will wait (potentially forever) for the medium-priority threads to complete.

One way to deal with this problem is to use a priority inheritance protocol. If a low-priority thread holds a mutex and a high-priority thread attempts to acquire the mutex, the priority of the thread holding the mutex is temporarily raised to the level of the thread waiting for the mutex. This allows the low-priority thread to keep running until it releases the mutex; as soon as it does, its priority is reduced back to its previous level and the high-priority thread acquires the mutex.

Ice supports the priority inheritance protocol on many POSIX platforms. (Windows does not provide such a protocol.)

For POSIX platforms that support the priority inheritance protocol, mutexes by default do not use it. You can use the `getDefaultMutexProtocol` function to retrieve the current default for your platform:

C++

```
namespace IceUtil {
    enum MutexProtocol { PrioInherit, PrioNone };

    MutexProtocol getDefaultMutexProtocol();
}
```

On Posix systems that do not support priority inheritance and on Windows, this function always returns `PrioNone`.

The return value of `getDefaultMutexProtocol` determines whether a default-constructed `Mutex` or `RecMutex` uses priority inheritance. By default, this function returns `PrioNone`. You can override this default by explicitly specifying a different protocol when you construct a `Mutex` or `RecMutex`. On Windows, if you specify `PrioInherit` when you construct a mutex, the setting is ignored and the mutex is constructed as if you had specified `PrioNone`.

To change the value returned by `getDefaultMutexProtocol`, you can edit `cpp/config/Make.rules`, modify the value of the `DEFAULT_MUTEX_PROTOCOL` macro, and then rebuild the `IceUtil` library.

See Also

- The C++ Mutex Class
- The C++ RecMutex Class

Copyright 2016, ZeroC, Inc.

Portable Signal Handling in C++

The `IceUtil::CtrlCHandler` class provides a portable mechanism to handle Ctrl+C and similar signals sent to a C++ process. On Windows, `IceUtil::CtrlCHandler` is a wrapper for `SetConsoleCtrlHandler`; on POSIX platforms, it handles `SIGHUP`, `SIGTERM` and `SIGINT` with a dedicated thread that waits for these signals using `sigwait`. Signals are handled by a callback function implemented and registered by the user. The callback is a simple function that takes an `int` (the signal number) and returns `void`; it must not throw any exception:

C++

```
namespace IceUtil {

    typedef void (*CtrlCHandlerCallback)(int);

    class CtrlCHandler {
    public:
        CtrlCHandler(CtrlCHandlerCallback = 0);
        ~CtrlCHandler();

        void setCallback(CtrlCHandlerCallback);
        CtrlCHandlerCallback getCallback() const;
    };
}
```

The member functions of `CtrlCHandler` behave as follows:

- `CtrlCHandler`
  Constructs an instance with a callback function. Only one instance of `CtrlCHandler` can exist in a process at a given moment in time. On POSIX platforms, the constructor masks `SIGHUP`, `SIGTERM` and `SIGINT`, then starts a thread that waits for these signals using `sigwait`. For signal masking to work properly, it is imperative that the `CtrlCHandler` instance be created before starting any thread, and in particular before initializing an Ice communicator.
- `~CtrlCHandler`
  Destroys the instance, after which the default signal processing behavior is restored on Windows (`TerminateProcess`). On POSIX platforms, the "sigwait" thread is terminated and joined, but the signal mask remains unchanged, so subsequent signals are ignored.
- `setCallback`
  Sets a new callback function.
- `getCallback`
  Gets the current callback function.

It is legal specify a value of zero (0) for the callback function, in which case signals are caught and ignored until a non-zero callback function is set.

A typical use for `CtrlCHandler` is to shutdown a communicator in an Ice server. For example, the `Ice::Application` class uses a `CtrlCHandler` in its implementation.

See Also

- The Server-Side main Function in C++

Copyright 2016, ZeroC, Inc.

The C++ AbstractMutex Class

`AbstractMutex` defines a mutex base interface used by the Freeze background save evictor. The interface allows the evictor to synchronize with servants that are stored in a Freeze database. The class has the following definition:

C++

```
class AbstractMutex {
public:
    typedef LockT<AbstractMutex> Lock;
    typedef TryLockT<AbstractMutex> TryLock;

    virtual ~AbstractMutex();

    virtual void lock() const = 0;
    virtual void unlock() const = 0;
    virtual bool tryLock() const = 0;
};
```

This class definition is provided in `IceUtil/AbstractMutex.h`. The same header file also defines a few template implementation classes that specialize `AbstractMutex`, as described below.

`AbstractMutexI`

This template class implements `AbstractMutex` by forwarding all member functions to its template argument:

C++

```
template <typename T>
class AbstractMutexI : public AbstractMutex, public T {
public:
    typedef LockT<AbstractMutexI> Lock;
    typedef TryLockT<AbstractMutexI> TryLock;

    virtual void lock() const {
        T::lock();
    }

    virtual void unlock() const {
        T::unlock();
    }

    virtual bool tryLock() const {
        return T::tryLock();
    }

    virtual ~AbstractMutexI() {}
};
```

`AbstractMutexReadI`

This template class implements a read lock by forwarding the `lock` and `tryLock` functions to the `readLock` and `tryReadLock` functions of its template argument:

Copyright 2016, ZeroC, Inc.

C++

```
template <typename T>
class AbstractMutexReadI : public AbstractMutex, public T {
public:
    typedef LockT<AbstractMutexReadI> Lock;
    typedef TryLockT<AbstractMutexReadI> TryLock;

    virtual void lock() const {
        T::readLock();
    }

    virtual void unlock() const {
        T::unlock();
    }

    virtual bool tryLock() const {
        return T::tryReadLock();
    }

    virtual ~AbstractMutexReadI() {}
};
```

`AbstractMutexWriteI`

This template class implements a write lock by forwarding the `lock` and `tryLock` functions to the `writeLock` and `tryWriteLock` function s of its template argument:

C++

```
template <typename T>
class AbstractMutexWriteI : public AbstractMutex, public T {
public:
    typedef LockT<AbstractMutexWriteI> Lock;
    typedef TryLockT<AbstractMutexWriteI> TryLock;

    virtual void lock() const {
        T::writeLock();
    }

    virtual void unlock() const {
        T::unlock();
    }

    virtual bool tryLock() const {
        return T::tryWriteLock();
    }

    virtual ~AbstractMutexWriteI() {}
};
```

Apart from use with Freeze servants, these templates are also useful if, for example, you want to implement your own evictor.

Copyright 2016, ZeroC, Inc.

See Also

- Background Save Evictor

Copyright 2016, ZeroC, Inc.

The C++ Cache Template

This class allows you to efficiently maintain a cache that is backed by secondary storage, such as a Berkeley DB database, without holding a lock on the entire cache while values are being loaded from the database. If you want to create evictors for servants that store their state in a database, the `Cache` class can simplify your evictor implementation considerably.

The `Cache` class has the following interface:

C++

```
template<typename Key, typename Value>
class Cache {
public:
    typedef typename std::map</* ... */, /* ... */>::iterator Position;

    bool pin(const Key& k, const Handle<Value>& v);
    Handle<Value> pin(const Key& k);
    void unpin(Position p);

    Handle<Value> putIfAbsent(const Key& k, const Handle<Value>& v);

    Handle<Value> getIfPinned(const Key&, bool = false) const;

    void clear();
    size_t size() const;


protected:
    virtual Handle<Value> load(const Key& k) = 0;
    virtual void pinned(const Handle<Value>& v, Position p);

    virtual ~Cache();
};
```

Note that `Cache` is an abstract base class — you must derive a concrete implementation from `Cache` and provide an implementation of the `load` and, optionally, of the `pinned` member function.

Internally, a `Cache` maintains a map of name-value pairs. The key and value type of the map are supplied by the `Key` and `Value` template arguments, respectively. The implementation of `Cache` takes care of maintaining the map; in particular, it ensures that concurrent lookups by callers are possible without blocking even if some of the callers are currently loading values from the backing store. In turn, this is useful for evictor implementations, such as the Freeze background save evictor. The `Cache` class does not limit the number of entries in the cache — it is the job of the evictor implementation to limit the map size by calling `unpin` on elements of the map that it wants to evict.

Your concrete implementation class must implement the `load` function, whose job it is to load the value for the key `k` from the backing store and to return a handle to that value. Note that `load` returns a value of type `IceUtil::Handle`, that is, the value must be heap-allocated and support the usual reference-counting functions for smart pointers. (The easiest way to achieve this is to derive the value from `IceUtil::Shared`.)

If `load` cannot locate a record for the given key because no such record exists, it must return a null handle. If `load` fails for some other reason, it can throw an exception, which is propagated back to the application code.

Your concrete implementation class typically will also override the `pinned` function (unless you want to have a cache that does not limit the number of entries; the provided default implementation of `pinned` is a no-op). The `Cache` implementation calls `pinned` whenever it has added a value to the map as a result of a call to `pin`; the `pinned` function is therefore a callback that allows the derived class to find out when a value has been added to the cache and informs the derived class of the value and its position in the cache.

The `Position` parameter is a `std::iterator` into the cache's internal map that records the position of the corresponding map entry. (Note that the element type of map is opaque, so you should not rely on knowledge of the cache's internal key and value types.) Your implementation of `pinned` must remember the position of the entry because that position is necessary to remove the corresponding entry from the cache again.

Copyright 2016, ZeroC, Inc.

The public member functions of `Cache` behave as follows:

```
bool pin(const Key& k, const Handle<Value>& v);
```

To add a key-value pair to the cache, your evictor can call `pin`. The return value is true if the key and value were added; a false return value indicates that the map already contained an entry with the given key and the original value for that key is unchanged.

`pin` calls `pinned` if it adds an entry.

This version of `pin` does *not* call `load` to retrieve the entry from backing store if it is not yet in the cache. This is useful when you add a newly-created object to the cache.

Once an entry is in the cache, it is guaranteed to remain in the cache at the same position in memory, and without its value being overwritten by another thread, until that entry is unpinned by a call to `unpin`.

```
Handle<Value> pin(const Key& k);
```

A second version of `pin` looks for the entry with the given key in the cache. If the entry is already in the cache, `pin` returns the entry's value. If no entry with the given key is in the cache, `pin` calls `load` to retrieve the corresponding entry. If `load` returns an entry, `pin` adds it to the cache and returns the entry's value. If the entry cannot be retrieved from the backing store, `pin` returns null.

`pin` calls `pinned` if it adds an entry.

The function is thread-safe, that is, it calls `load` only once all other threads have unpinned the entry.

Once an entry is in the cache, it is guaranteed to remain in the cache at the same position in memory, and without its value being overwritten by another thread, until that entry is unpinned by a call to `unpin`.

```
Handle<Value> putIfAbsent(const Key& k, const Handle<Value>& v);
```

This function adds a key-value pair to the cache and returns a smart pointer to the value. If the map already contains an entry with the given key, that entry's value remains unchanged and `putIfAbsent` returns its value. If no entry with the given key is in the cache, `putIfAbsent` calls `load` to retrieve the corresponding entry. If `load` returns an entry, `putIfAbsent` adds it to the cache and returns the entry's value. If the entry cannot be retrieved from the backing store, `putIfAbsent` returns null.

`putIfAbsent` calls `pinned` if it adds an entry.

The function is thread-safe, that is, it calls `load` only once all other threads have unpinned the entry.

Once an entry is in the cache, it is guaranteed to remain in the cache at the same position in memory, and without its value being overwritten by another thread, until that entry is unpinned by a call to `unpin`.

```
Handle<Value> getIfPinned(const Key& k, bool wait = false) const;
```

This function returns the value stored for the key `k`.

- If an entry for the given key is in the map, the function returns the value immediately, regardless of the value of `wait`.
- If no entry for the given key is in the map and the `wait` parameter is false, the function returns a null handle.
- If no entry for the given key is in the map and the `wait` parameter is true, the function blocks the calling thread if another thread is currently attempting to load the same entry; once the other thread completes, `getIfPinned` completes and returns the value added by the other thread.

```
void unpin(Position p);
```

This function removes an entry from the map. The iterator `p` determines which entry to remove. (It must be an iterator that previously was passed to `pinned`.) The iterator `p` is invalidated by this operation, so you must not use it again once `unpin` returns. (Note that the `Cache` implementation ensures that updates to the map never invalidate iterators to existing entries in the map; `unpin` invalidates only the iterator for the removed entry.)

```
void clear();
```

This function removes all entries in the map.

```
size_t size() const;
```

This function returns the number of entries in the map.

See Also

- Servant Evictors
- The C++ Handle Template
- The C++ Shared and SimpleShared Classes

Copyright 2016, ZeroC, Inc.

The C++ Exception Class

This class is at the root of the derivation tree for Ice exceptions and encapsulates functionality that is common to all `Ice` and `IceUtil` exceptions:

C++

```
class Exception : public std::exception {
public:
    Exception();
    Exception(const char* file, int line);
    virtual ~Exception() throw();

    virtual std::string ice_name() const;
    virtual void ice_print(std::ostream&) const;
    virtual const char* what() const throw();
    virtual Exception* ice_clone() const;
    virtual void ice_throw() const;
    const char* ice_file() const;
    int ice_line() const;
    const std::string& ice_stackTrace() const;
};
```

The second constructor stores a file name and line number in the exception that are returned by the `ice_file` and `ice_line` member functions, respectively. This allows you to identify the source of an exception by passing the `__FILE__` and `__LINE__` preprocessor macros to the constructor.

The `what` member function is a synonym for `ice_print`. The default implementation of `ice_print` prints the file name, line number, and the name of the exception.

The `ice_stackTrace` function returns the full stack trace when the exception was constructed, or an empty string, depending on the value of the `Ice.PrintStackTraces` property.

The remaining member functions are described in the C++ Mapping for Exceptions.

See Also

- C++ Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

The C++ generateUUID Function

Universally-unique identifiers (UUIDs) are often used in the identities of Ice objects. The C++ standard does not include a function for generating UUIDs, therefore Ice provides the `IceUtil::generateUUID` function for use in portable applications. The signature of `generateUUID` is:

C++

```
std::string generateUUID();
```

The function returns a string like the following:

```
02b066f5-c762-431c-8dd3-9b1941355e41
```

Each invocation returns a new identifier that differs from all previous ones.
See Also

- Object Identity

Copyright 2016, ZeroC, Inc.

The C++ Handle Template

`IceUtil::Handle` implements a smart reference-counted pointer type. Smart pointers are used to guarantee automatic deletion of heap-allocated class instances.

`Handle` is a template class with the following interface:

C++

```
template<typename T>
class Handle : /* ... */ {
public:

    typedef T element_type;

    T* _ptr;

    T* operator->() const;
    T& operator*() const;
    T* get() const;

    operator bool() const;

    void swap(HandleBase& other);

    Handle(T* p = 0);

    template<typename Y>
    Handle(const Handle<Y>& r);

    Handle(const Handle& r);

    ~Handle();

    Handle& operator=(T* p);

    template<typename Y>
    Handle& operator=(const Handle<Y>& r);

    Handle& operator=(const Handle& r);

    template<class Y>
    static Handle dynamicCast(const HandleBase<Y>& r);

    template<class Y>
    static Handle dynamicCast(Y* p);
};

template<typename T, typename U>
bool operator==(const Handle<T>& lhs, const Handle<U>& rhs);

template<typename T, typename U>
bool operator!=(const Handle<T>& lhs, const Handle<U>& rhs);

template<typename T, typename U>
```


Copyright 2016, ZeroC, Inc.

```
bool operator<(const Handle<T>& lhs, const Handle<U>& rhs);

template<typename T, typename U>
bool operator<=(const Handle<T>& lhs, const Handle<U>& rhs);

template<typename T, typename U>
bool operator>(const Handle<T>& lhs, const Handle<U>& rhs);
```


Copyright 2016, ZeroC, Inc.

```
template<typename T, typename U>
bool operator>=(const Handle<T>& lhs, const Handle<U>& rhs);
```

Note that the actual implementation is split into a base and a derived class. For simplicity, we show the combined interface here. If you want to see the full implementation detail, it can be found in `IceUtil/Handle.h`.

The template argument must be a class that derives from `Shared` or `SimpleShared` (or that implements reference counting with the same interface as these classes).

This is quite a large interface, but all it really does is to faithfully mimic the behavior of ordinary C++ class instance pointers. Rather than discussing each member function in detail, we provide a simple overview here that outlines the most important points. Please see the discussion of Ice objects for more examples of using smart pointers.

`element_type`

This type definition follows the STL convention of defining the element type with the fixed name `element_type` so you can use it for template programming or the definition of generic containers.

`_ptr`

This data member stores the pointer to the underlying heap-allocated class instance.

Constructors, copy constructor, and assignment operators

These member functions allow you to construct, copy, and assign smart pointers as if they were ordinary pointers. In particular, the constructor and assignment operator are overloaded to work with raw C++ class instance pointers, which results in the "adoption" of the raw pointer by the smart pointer. For example, the following code works correctly and does not cause a memory leak:

C++

```
typedef Handle<MyClass> MyClassPtr;

void foo(const MyClassPtr&);

// ...

foo(new MyClass); // OK, no leak here.
```

`operator->`, `operator*`, and `get`

The arrow and indirection operators allow you to apply the usual pointer syntax to smart pointers to use the target of a smart pointer. The `get` member function returns the class instance pointer to the underlying reference-counted class instance; the return value is the value of `_ptr`.

`dynamicCast`

This member function works exactly like a C++ `dynamic_cast`: it tests whether the argument supports the specified type and, if so, returns a non-null pointer; if the target does not support the specified type, it returns null.

The reason for not using an actual `dynamic_cast` and using a `dynamicCast` function instead is that `dynamic_cast` only operates on pointer types, but `IceUtil::Handle` is a class.

For example:

Copyright 2016, ZeroC, Inc.

C++

```
MyClassPtr p = ...;
MyOtherClassPtr o = ...;

o = MyOtherClassPtr::dynamicCast(p);
if (o)
{
    // o points at an instance of type MyOtherClass.
}
else
{
    // p points at something that is
    // not compatible with MyOtherClass.
}
```

Note that this example also illustrates the use of `operator bool`: when used in a boolean context, a smart pointer returns true if it is non-null and false otherwise.
Comparison operators: `==, !=, <, <=, >, >=`

The comparison operators for smart pointers delegate to the operators of the underlying class, therefore the author of the reference-counted class defines the semantics of smart pointer comparison. For example, in the case of Slice classes, the base class `Ice::Object` implements the comparison operators in terms of pointer addresses. On the other hand, the base class for Slice proxies implements comparison using value semantics.
See Also

- Smart Pointers for Classes
- C++ Mapping for Interfaces
- The C++ Shared and SimpleShared Classes

Copyright 2016, ZeroC, Inc.

The C++ Handle Template Adaptors

`IceUtil` provides adaptors that support use of smart pointers with STL algorithms. Each template function returns a corresponding function object that is for use by an STL algorithm. The adaptors are defined in the header `IceUtil/Functional.h`.

Here is a list of the adaptors:

```
memFun
memFun1
voidMemFun
voidMemFun1

secondMemFun
secondMemFun1
secondVoidMemFun
secondVoidMemFun1

constMemFun
constMemFun1
constVoidMemFun
constVoidMemFun1

secondConstMemFun
secondConstMemFun1
secondConstVoidMemFun
secondConstVoidMemFun1
```

As you can see, the adaptors are in two groups. The first group operates on non-const smart pointers, whereas the second group operates on `const` smart pointers (for example, on smart pointers declared as `const MyClassPtr`).

Each group is further divided into two sub-groups. The adaptors in the first group operate on the target of a smart pointer, whereas the `second<`*name*`>` adapters operate on the second element of a pair, where that element is a smart pointer.

Each of the four sub-groups contains four adaptors:
`memFun`

This adaptor is used for member functions that return a value and do not accept an argument. For example:

Copyright 2016, ZeroC, Inc.

C++

```
class MyClass : public IceUtil::Shared {
public:
    MyClass(int i) : _i(i) {}
    int getVal() { return _i; }
private:
    int _i;
};

typedef IceUtil::Handle<MyClass> MyClassPtr;

// ...

vector<MyClassPtr> mcp;
mcp.push_back(new MyClass(42));
mcp.push_back(new MyClass(99));

transform(mcp.begin(), mcp.end(),
          ostream_iterator<int>(cout, " "),
          IceUtil::memFun(&MyClass::getVal));
cout << endl;
```

This code invokes the member function `getVal` on each instance that is pointed at by smart pointers in the vector `mcp` and prints the return value of `getVal` on `cout`, separated by spaces. The output from this code is:

```
42 99
```

`memFun1`

This adaptor is used for member functions that return a value and accept a single argument. For example:

Copyright 2016, ZeroC, Inc.

C++

```
class MyClass : public IceUtil::Shared {
public:
    MyClass(int i) : _i(i) {}
    int plus(int v) { return _i + v; }
private:
    int _i;
};

typedef IceUtil::Handle<MyClass> MyClassPtr;

// ...

vector<MyClassPtr> mcp;
mcp.push_back(new MyClass(2));
mcp.push_back(new MyClass(4));
mcp.push_back(new MyClass(6));

int A[3] = { 5, 7, 9 };
transform(mcp.begin(), mcp.end(), A,
          ostream_iterator<int>(cout, " "),
          IceUtil::memFun1(&MyClass::plus));
cout << endl;
```

This code invokes the member function `plus` on each instance that is pointed at by smart pointers in the vector `mcp` and prints the return value of a call to `plus` on `cout`, separated by spaces. The calls to `plus` are successively passed the values stored in the array `A`. The output from this code is:

```
7 11 15
```

`voidMemFun`

This adaptor is used for member functions that do not return a value and do not accept an argument. For example:

Copyright 2016, ZeroC, Inc.

C++

```
class MyClass : public IceUtil::Shared {
public:
    MyClass(int i) : _i(i) {}
    void print() { cout << _i << endl; }
private:
    int _i;
};

typedef IceUtil::Handle<MyClass> MyClassPtr;

// ...

vector<MyClassPtr> mcp;
mcp.push_back(new MyClass(2));
mcp.push_back(new MyClass(4));
mcp.push_back(new MyClass(6));

for_each(mcp.begin(), mcp.end(), IceUtil::voidMemFun(&MyClass::print));
```

This code invokes the member function `print` on each instance that is pointed at by smart pointers in the vector `mcp`. The output from this code is:

```
2
4
6
```

`voidMemFun1`

This adaptor is used for member functions that do not return a value and accept a single argument. For example:

Copyright 2016, ZeroC, Inc.

C++

```
class MyClass : public IceUtil::Shared {
public:
    MyClass(int i) : _i(i) {}
    void printPlus(int v) { cout << _i + v << endl; }
private:
    int _i;
};

typedef IceUtil::Handle<MyClass> MyClassPtr;

vector<MyClassPtr> mcp;
mcp.push_back(new MyClass(2));
mcp.push_back(new MyClass(4));
mcp.push_back(new MyClass(6));

for_each(
    mcp.begin(), mcp.end(),
    bind2nd(IceUtil::voidMemFun1(&MyClass::printPlus), 3));
```

This code invokes the member function `printPlus` on each instance that is pointed at by smart pointers in the vector `mcp`. The output from this code is:

```
5
7
9
```

As mentioned earlier, the `second<`*name*`>` versions of the adaptors operate on the second element of a `std::pair<T1, T2>`, where `T2` must be a smart pointer. Most commonly, these adaptors are used to apply an algorithm to each lookup value of a map or multi-map. Here is an example:

Copyright 2016, ZeroC, Inc.

C++

```
class MyClass : public IceUtil::Shared {
public:
    MyClass(int i) : _i(i) {}
    int plus(int v) { return _i + v; }
private:
    int _i;
};

typedef IceUtil::Handle<MyClass> MyClassPtr;

// ...

map<string, MyClassPtr> m;
m["two"] = new MyClass(2);
m["four"] = new MyClass(4);
m["six"] = new MyClass(6);

int A[3] = { 5, 7, 9 };
transform(
    m.begin(), m.end(), A,
    ostream_iterator<int>(cout, " "),
    IceUtil::secondMemFun1<int, string, MyClass>(&MyClass::plus));
```

This code invokes the `plus` member function on the class instance denoted by the `second` smart pointer member of each pair in the dictionary `m`. The output from this code is:

```
9 13 11
```

Note that `secondMemFun1` is a template that requires three arguments: the return type of the member function to be invoked, the key type of the dictionary, and the type of the class that is pointed at by the smart pointer.

In general, the `second`<*name*> adaptors require the following template arguments:

C++

```
secondMemFun<R, K, T>
secondMemFun1<R, K, T>
secondVoidMemFun<K, T>
secondVoidMemFun<K, T>
```

where `R` is the return type of the member function, `K` is the type of the first member of the pair, and `T` is the class that contains the member function.

See Also

- The C++ Handle Template

Copyright 2016, ZeroC, Inc.

The C++ ScopedArray Template

`IceUtil::ScopedArray` is a smart pointer class similar to `Handle`. However, instead of managing the memory for class instances, `ScopedArray` manages memory for an array. This class is provided mainly for use with the stream API. However, you can use it with arrays for other purposes.

Here is the definition of the template in full:

C++

```
template<typename T>
class ScopedArray : private IceUtil::noncopyable
{
public:
    explicit ScopedArray(T* ptr = 0)
        : _ptr(ptr) { }

    ScopedArray(const ScopedArray& other) {
        _ptr = other._ptr;
        const_cast<ScopedArray&>(other)._ptr = 0;
    }

    ~ScopedArray() {
        if (_ptr != 0)
            delete[] _ptr;
    }

    void reset(T* ptr = 0) {
        assert(ptr == 0 || ptr != _ptr);
        if (_ptr != 0)
            delete[] _ptr;
        _ptr = ptr;
    }

    T& operator[](size_t i) const {
        assert(_ptr != 0);
        assert(i >= 0);
        return _ptr[i];
    }

    T* get() const {
        return _ptr;
    }

    void swap(ScopedArray& a) {
        T* tmp = a._ptr;
        a._ptr = _ptr;
        _ptr = tmp;
    }

private:
    T* _ptr;
};
```


Copyright 2016, ZeroC, Inc.

The class allows you to allocate an array on the heap and assign its pointer to a `ScopedArray` instance. When the instance goes out of scope, it calls `delete[]` on the array, so you do not need to deallocate the array explicitly yourself. This greatly reduces the risk of a memory leak due to an early return or uncaught exception.

See Also

- C++ Streaming Interfaces

Copyright 2016, ZeroC, Inc.

The C++ Shared and SimpleShared Classes

`IceUtil::Shared` and `IceUtil::SimpleShared` are base classes that implement the reference-counting mechanism for smart pointers. The two classes provide identical interfaces; the difference between `Shared` and `SimpleShared` is that `SimpleShared` is not thread-safe and, therefore, can only be used if the corresponding class instances are accessed only by a single thread. (`SimpleShared` is marginally faster than `Shared` because it avoids the locking overhead that is incurred by `Shared`.)

The interface of `Shared` looks as follows. (Because `SimpleShared` has the same interface, we do not show it separately here.)

C++

```
class Shared {
public:
    Shared();
    Shared(const Shared&);
    virtual ~Shared();

    Shared& operator=(const Shared&);

    virtual void __incRef();
    virtual void __decRef();
    virtual int __getRef() const;
    virtual void __setNoDelete(bool);
};
```

The class maintains a reference that is initialized to zero by the constructor. `__incRef` increments the reference count and `__decRef` decrements it. If, during a call to `__decRef`, after decrementing the reference count, the reference count drops to zero, `__decRef` calls `delete this`, which causes the corresponding class instance to delete itself. The copy constructor increments the reference count of the copied instance, and the assignment operator increments the reference count of the source and decrements the reference count of the target.

The `__getRef` member function returns the value of the reference count and is useful mainly for debugging.

The `__setNoDelete` member function can be used to temporarily disable self-deletion and re-enable it again. This provides exception safety when you initialize a smart pointer with the `this` pointer of a class instance during construction.

To create a class that is reference-counted, you simply derive the class from `Shared` and define a smart pointer type for the class, for example:

C++

```
class MyClass : public IceUtil::Shared {
    // ...
};

typedef IceUtil::Handle<MyClass> MyClassPtr;
```

See Also

- The C++ Handle Template
- Smart Pointers for Classes

Copyright 2016, ZeroC, Inc.

The C++ Time Class

The `Time` class provides basic facilities for getting the current time, constructing time intervals, adding and subtracting times, and comparing times:

C++

```
namespace IceUtil {

    typedef ... Int64;

    class Time {
    public:
        enum Clock { Realtime, Monotonic };
        Time(Clock = Realtime);
        static Time now();
        static Time seconds(Int64);
        static Time milliSeconds(Int64);
        static Time microSeconds(Int64);

        Int64 toSeconds() const;
        Int64 toMilliSeconds() const;
        Int64 toMicroSeconds() const;

        double toSecondsDouble() const;
        double toMilliSecondsDouble() const;
        double toMicroSecondsDouble() const;

        std::string toDateTime() const;
        std::string toDuration() const;

        Time operator-() const;

        Time operator-(const Time&) const;
        Time operator+(const Time&) const;

        Time operator*(int) const;
        Time operator*(Int64) const;
        Time operator*(double) const;

        double operator/(const Time&) const;
        Time operator/(int) const;
        Time operator/(Int64) const;
        Time operator/(double) const;

        Time& operator-=(const Time&);
        Time& operator+=(const Time&);

        Time& operator*=(int);
        Time& operator*=(Int64);
        Time& operator*=(double);

        Time& operator/=(int);
        Time& operator/=(Int64);
```


Copyright 2016, ZeroC, Inc.

```
        Time& operator/=(double);

        bool operator<(const Time&) const;
        bool operator<=(const Time&) const;
        bool operator>(const Time&) const;
        bool operator>=(const Time&) const;
        bool operator==(const Time&) const;
        bool operator!=(const Time&) const;

#ifndef _WIN32
        operator timeval() const;
#endif
    };
```


Copyright 2016, ZeroC, Inc.

```
    std::ostream& operator<<(std::ostream&, const Time&);
}
```

The member functions behave as follows:
`Time`

Internally, the `Time` class stores ticks in microsecond units. For absolute time, this is the number of microseconds since the Unix epoch (00:00:00 UTC on 1 Jan. 1970). For durations, this is the number of microseconds in the duration. The default constructor initializes the tick count to zero and selects the real-time clock. Constructing `Time` with an argument of `Monotonic` selects the monotonic clock on platforms that support it; the real-time clock is used on other platforms.
`now`

This function constructs a `Time` object that is initialized to the current time of day.
`seconds, milliSeconds, microSeconds`

These functions construct `Time` objects from the argument in the specified units. For example, the following statement creates a time duration of one minute:

C++

```
IceUtil::Time t = IceUtil::Time::seconds(60);
```

`toSeconds, toMilliSeconds, toMicroSeconds`

The member functions provide explicit conversion of a duration to seconds, milliseconds, and microseconds, respectively. The return value is a 64-bit signed integer (`IceUtil::Int64`). For example:

C++

```
IceUtil::Time t = IceUtil::Time::milliSeconds(2000);
IceUtil::Int64 secs = t.toSeconds(); // Returns 2
```

`toSecondsDouble, toMilliSecondsDouble, toMicroSecondsDouble`

The member functions provide explicit conversion of a duration to seconds, milliseconds, and microseconds, respectively. The return value is of type `double`.
`toDateTime`

This function returns a human-readable representation of a `Time` value as a date and time.
`toDuration`

This function returns a human-readable representation of a `Time` value as a duration.
`Operators`

`Time` provides operators that allow you to add, subtract, multiply, and divide times. For example:

C++

```
IceUtil::Time oneMinute = IceUtil::Time::seconds(60);
IceUtil::Time oneMinuteAgo = IceUtil::Time::now() - oneMinute;
```


Copyright 2016, ZeroC, Inc.

The multiplication and division operators permit you to multiply and divide a duration. Note that these operators provide overloads for `int`, `long long`, and `double`.

The comparison operators allow you to compare times and time intervals with each other, for example:

```
IceUtil::Time oneMinute  = IceUtil::Time::seconds(60);
IceUtil::Time twoMinutes = IceUtil::Time::seconds(120);
assert(oneMinute < twoMinutes);
```

The `timeval` operator converts a `Time` object to a `struct timeval`, defined as follows:

C++

```
struct timeval {
    long tv_sec;
    long tv_usec;
};
```

The conversion is useful for API calls that require a `struct timeval` argument, such as `select`. To convert a duration into a `timeval` structure, simply assign a `Time` object to a `struct timeval`:

C++

```
IceUtil::Time oneMinute = IceUtil::Time::seconds(60);
struct timeval tv;
tv = t;
```

Note that this member function is not available under Windows.
`std::ostream& operator<<(std::ostream&, Time&);`

This operator prints the number of whole seconds since the epoch.
See Also

- The C++ Timer and TimerTask Classes

Copyright 2016, ZeroC, Inc.

The C++ Timer and TimerTask Classes

The `Timer` class allows you to schedule some code for once-only or repeated execution after some time interval elapses. The code to be executed resides in a class you derive from `TimerTask`:

C++

```
class Timer;
typedef IceUtil::Handle<Timer> TimerPtr;

class TimerTask : virtual public IceUtil::Shared {
public:
    virtual ~TimerTask() { }
    virtual void runTimerTask() = 0;
};

typedef IceUtil::Handle<TimerTask> TimerTaskPtr;
```

Your derived class must override the `runTimerTask` member function; the code in this method is executed by the timer. If the code you want to run requires access to some program state, you can pass that state into the constructor of your class or, alternatively, set that state via member functions of your class before scheduling it with a timer.

The `Timer` class invokes the `runTimerTask` member function to run your code. The class has the following definition:

C++

```
class Timer : /* ... */ {
public:
    Timer();
    Timer(int priority);

    void schedule(const TimerTaskPtr& task, const IceUtil::Time& interval);

    void scheduleRepeated(const TimerTaskPtr& task, const IceUtil::Time& interval);

    bool cancel(const TimerTaskPtr& task);

    void destroy();
};

typedef IceUtil::Handle<Timer> TimerPtr;
```

Intervals are specified using `Time` objects.

The constructor is overloaded to allow you specify a thread priority. The priority controls the priority of the thread that executes your task.

The `schedule` member function schedules an instance of your timer task for once-only execution after the specified time interval has elapsed. Your code is executed by a separate thread that is created by the `Timer` class. The function throws an `IllegalArgumentException` if you invoke it on a destroyed timer.

The `scheduleRepeated` member function runs your task repeatedly, at the specified time interval. Your code is executed by a separate thread that is created by the `Timer` class; the same thread is used every time your code runs. The function throws an `IllegalArgumentException` if you invoke it on a destroyed timer.

If your code throws an exception, the `Timer` class ignores the exception, that is, for a task that is scheduled to run repeatedly, an exception in the current execution does not cancel the next execution.

Copyright 2016, ZeroC, Inc.

If your code takes longer to execute than the time interval you have specified for repeated execution, the second execution is delayed accordingly. For example, if you ask for repeated execution once every five seconds, and your code takes ten seconds to complete, then the second execution of your task starts five seconds after the previous execution finishes, that is, the interval specifies the wait time between successive executions.

A `TimerTask` instance that has already been scheduled with a `Timer` instance cannot be scheduled again with the same `Timer` instance until the task has completed or been canceled.

For a single `Timer` instance, the execution of all registered tasks is serialized. The wait interval applies on a per-task basis so, if you schedule task A at an interval of five seconds, and task B at an interval of ten seconds, successive runs of task A start no sooner than five seconds after the previous task A has finished, and successive runs of task B start no sooner than ten seconds after the previous task B has finished. If, at the time a task is scheduled to run, another task is still running, the new task's execution is delayed until the previous task has finished.

If you want scheduled tasks to run concurrently, you can create several `Timer` instances; tasks then execute in as many threads concurrently as there are `Timer` instances.

The `cancel` member function removes a task from a timer's schedule. In other words, it stops a task that is scheduled from being executed. If you cancel a task while it is executing, `cancel` returns immediately and the currently running task is allowed to complete normally; that is, `cancel` does not wait for any currently running task to complete.

The return value is true if `cancel` removed the task from the schedule. This is the case if you invoke `cancel` on a task that is scheduled for repeated execution and this was the first time you cancelled that task; subsequent calls to `cancel` return false. Calling `cancel` on a task scheduled for once-only execution always returns false, as does calling `cancel` on a destroyed timer.

The `destroy` member function removes all tasks from the timer's schedule. If you call `destroy` from any thread other than the timer's own execution thread, it joins with the currently executing task (if any), so the function does not return until the current task has completed. If you call `destroy` from the timer's own execution thread, it instead detaches the timer's execution thread. Calling `destroy` a second time on the same `Timer` instance has no effect. Similarly, calling `cancel` on a destroyed timer has no effect.

Note that you must call `destroy` on a `Timer` instance before allowing it to go out of scope; failing to do so causes undefined behavior.

Calls to `schedule` or `scheduleRepeated` on a destroyed timer raises an `IceUtil::IllegalArgumentException`.

See Also

- The C++ Time Class
- The C++ Thread Classes

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping

Topics

- Client-Side Slice-to-C-Sharp Mapping
- Server-Side Slice-to-C-Sharp Mapping
- The .NET Utility Library

Copyright 2016, ZeroC, Inc.

Client-Side Slice-to-C-Sharp Mapping

The client-side Slice-to-C# mapping defines how Slice data types are translated to C# types, and how clients invoke operations, pass parameters, and handle errors. Much of the C# mapping is intuitive. For example, by default, Slice sequences map to C# arrays, so there is little you have learn in order to use Slice sequences in C#.

The C# API to the Ice run time is fully thread-safe. Obviously, you must still synchronize access to data from different threads. For example, if you have two threads sharing a sequence, you cannot safely have one thread insert into the sequence while another thread is iterating over the sequence. However, you only need to concern yourself with concurrent access to your own data — the Ice run time itself is fully thread safe, and none of the Ice API calls require you to acquire or release a lock before you safely can make the call.

Much of what appears in this chapter is reference material. We suggest that you skim the material on the initial reading and refer back to specific sections as needed. However, we recommend that you read at least the mappings for exceptions, interfaces, and operations in detail because these sections cover how to call operations from a client, pass parameters, and handle exceptions.

In order to use the C# mapping, you should need no more than the Slice definition of your application and knowledge of the C# mapping rules. In particular, looking through the generated code in order to discern how to use the C# mapping is likely to be inefficient, due to the amount of detail. Of course, occasionally, you may want to refer to the generated code to confirm a detail of the mapping, but we recommend that you otherwise use the material presented here to see how to write your client-side code.

The `Ice` Namespace

All of the APIs for the Ice run time are nested in the `Ice` namespace, to avoid clashes with definitions for other libraries or applications. Some of the contents of the `Ice` namespace are generated from Slice definitions; other parts of the `Ice` namespace provide special-purpose definitions that do not have a corresponding Slice definition. We will incrementally cover the contents of the `Ice` namespace throughout the remainder of the manual.

Topics

- C-Sharp Mapping for Identifiers
- C-Sharp Mapping for Modules
- C-Sharp Mapping for Built-In Types
- C-Sharp Mapping for Enumerations
- C-Sharp Mapping for Structures
- C-Sharp Mapping for Sequences
- C-Sharp Mapping for Dictionaries
- C-Sharp Collection Comparison
- C-Sharp Mapping for Constants
- C-Sharp Mapping for Exceptions
- C-Sharp Mapping for Interfaces
- C-Sharp Mapping for Operations
- C-Sharp Mapping for Classes
- C-Sharp Mapping for Optional Values
- Serializable Objects in C-Sharp
- C-Sharp Specific Metadata Directives
- Asynchronous Method Invocation (AMI) in C-Sharp
- slice2cs Command-Line Options
- Using Slice Checksums in C-Sharp
- Example of a File System Client in C-Sharp

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Identifiers

Slice identifiers map to an identical C# identifier. For example, the Slice identifier `Clock` becomes the C# identifier `Clock`. If a Slice identifier is the same as a C# keyword, the corresponding C# identifier is a *verbatim identifier* (an identifier prefixed with @). For example, the Slice identifier `while` is mapped as `@while`.

You should try to avoid such identifiers as much as possible.

The Slice-to-C# compiler generates classes that inherit from interfaces or base classes in the .NET framework. These interfaces and classes introduce a number of methods into derived classes. To avoid name clashes between Slice identifiers that happen to be the same as an inherited method, such identifiers are prefixed with `ice_` and suffixed with _ in the generated code. For example, the Slice identifier `Clone` maps to the C# identifier `ice_Clone_` if it would clash with an inherited `Clone`. The complete list of identifiers that are so changed is:

`Clone`	`Equals`	`Finalize`
`GetBaseException`	`GetHashCode`	`GetObjectData`
`GetType`	`MemberwiseClone`	`ReferenceEquals`
`ToString`	`checkedCast`	`uncheckedCast`

Note that Slice identifiers in this list are translated to the corresponding C# identifier only where necessary. For example, structures do not derive from `ICloneable`, so if a Slice structure contains a member named `Clone`, the corresponding C# structure's member is named `Clone` as well. On the other hand, classes do derive from `ICloneable`, so, if a Slice class contains a member named `Clone`, the corresponding C# class's member is named `ice_Clone_`.

Also note that, for the purpose of prefixing, Slice identifiers are case-insensitive, that is, both `Clone` and `clone` are escaped and map to `ice_Clone_` and `ice_clone_`, respectively.

See Also

- Identifiers That Are Keywords
- C-Sharp Mapping for Modules
- C-Sharp Mapping for Built-In Types
- C-Sharp Mapping for Enumerations
- C-Sharp Mapping for Structures
- C-Sharp Mapping for Sequences
- C-Sharp Mapping for Dictionaries
- C-Sharp Collection Comparison
- C-Sharp Mapping for Constants
- C-Sharp Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Modules

Slice modules map to C# namespaces with the same name as the Slice module. The mapping preserves the nesting of the Slice definitions. For example:

Slice

```
module M1 {
    // Definitions for M1 here...
    module M2 {
        // Definitions for M2 here...
    };
};

// ...

module M1 {     // Reopen M1
    // More definitions for M1 here...
};
```

This definition maps to the corresponding C# definitions:

C#

```
namespace M1
{
    namespace M2
    {
        // ...
    }
    // ...
}

// ...

namespace M1    // Reopen M1
{
    // ...
}
```

If a Slice module is reopened, the corresponding C# namespace is reopened as well.
See Also

- C-Sharp Mapping for Identifiers
- C-Sharp Mapping for Built-In Types
- C-Sharp Mapping for Enumerations
- C-Sharp Mapping for Structures
- C-Sharp Mapping for Sequences
- C-Sharp Mapping for Dictionaries
- C-Sharp Collection Comparison
- C-Sharp Mapping for Constants
- C-Sharp Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Built-In Types

The Slice built-in types are mapped to C# types as shown below:

Slice	C#
`bool`	`bool`
`byte`	`byte`
`short`	`short`
`int`	`int`
`long`	`long`
`float`	`float`
`double`	`double`
`string`	`string`

Mapping of Slice built-in types to C#.

See Also

- C-Sharp Mapping for Identifiers
- C-Sharp Mapping for Modules
- C-Sharp Mapping for Enumerations
- C-Sharp Mapping for Structures
- C-Sharp Mapping for Sequences
- C-Sharp Mapping for Dictionaries
- C-Sharp Collection Comparison
- C-Sharp Mapping for Constants
- C-Sharp Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Enumerations

A Slice enumeration maps to the corresponding enumeration in C#. For example:

Slice

```
enum Fruit { Apple, Pear, Orange };
```

Not surprisingly, the generated C# definition is very similar:

C#

```
public enum Fruit { Apple, Pear, Orange }
```

Suppose we modify the Slice definition to include a custom enumerator value:

Slice

```
enum Fruit { Apple, Pear = 3, Orange };
```

The generated C# definition now includes an explicit initializer for every enumerator:

C#

```
public enum Fruit { Apple = 0, Pear = 3, Orange = 4 }
```

See Also

- C-Sharp Mapping for Identifiers
- C-Sharp Mapping for Modules
- C-Sharp Mapping for Built-In Types
- C-Sharp Mapping for Structures
- C-Sharp Mapping for Sequences
- C-Sharp Mapping for Dictionaries
- C-Sharp Collection Comparison
- C-Sharp Mapping for Constants
- C-Sharp Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Structures

Ice for .NET supports two different mappings for Slice structures. By default, Slice structures map to C# structures if they (recursively) contain only value types. If a Slice structure (recursively) contains a string, proxy, class, sequence, or dictionary member, it maps to a C# class. A metadata directive allows you to force the mapping to a C# class for Slice structures that contain only value types.

In addition, for either mapping, you can control whether Slice data members are mapped to fields or to properties.

On this page:

- Structure Mapping for Structures in C#
- Class Mapping for Structures in C#
- Property Mapping for Structures in C#

Structure Mapping for Structures in C#

Consider the following structure:

Slice

```
struct Point {
    double x;
    double y;
};
```

This structure consists of only value types and so, by default, maps to a C# partial structure:

C#

```
public partial struct Point
{
    public double x;
    public double y;

    public Point(double x, double y);

    public override int GetHashCode();
    public override bool Equals(object other);

    public static bool operator==(Point lhs, Point rhs);
    public static bool operator!=(Point lhs, Point rhs);
}
```

For each data member in the Slice definition, the C# structure contains a corresponding public data member of the same name.

The generated constructor accepts one argument for each structure member, in the order in which they are defined in the Slice definition. This allows you to construct and initialize a structure in a single statement:

C#

```
Point p = new Point(5.1, 7.8);
```


Copyright 2016, ZeroC, Inc.

Note that C# does not allow a value type to declare a default constructor or to assign default values to data members.

The structure overrides the `GetHashCode` and `Equals` methods to allow you to use it as the key type of a dictionary. (Note that the static two-argument version of `Equals` is inherited from `System.Object`.) Two structures are equal if (recursively) all their data members are equal. Otherwise, they are not equal. For structures that contain reference types, `Equals` performs a deep comparison; that is, reference types are compared for value equality, not reference equality.

Class Mapping for Structures in C#

The mapping for Slice structures to C# structures provides value semantics. Usually, this is appropriate, but there are situations where you may want to change this:

- If you use structures as members of a collection, each access to an element of the collection incurs the cost of boxing or unboxing. Depending on your situation, the performance penalty may be noticeable.
- On occasion, it is useful to be able to assign null to a structure, for example, to support "not there" semantics (such as when implementing parameters that are conceptually optional).

To allow you to choose the correct performance and functionality trade-off, the Slice-to-C# compiler provides an alternative mapping of structures to classes, for example:

Slice

```
["clr:class"] struct Point {
    double x;
    double y;
};
```

The `"clr:class"` metadata directive instructs the Slice-to-C# compiler to generate a mapping to a C# partial class for this structure. The generated code is almost identical, except that the keyword `struct` is replaced by the keyword `class` and that the class has a default constructor and inherits from `ICloneable`:

C#

```
public partial class Point : _System.ICloneable
{
    public double x;
    public double y;

    public Point();
    public Point(double x, double y);

    public object Clone();

    public override int GetHashCode();
    public override bool Equals(object other);

    public static bool operator==(Point lhs, Point rhs);
    public static bool operator!=(Point lhs, Point rhs);
}
```

Some of the generated marshaling code differs for the class mapping of structures, but this is irrelevant to application code.

Copyright 2016, ZeroC, Inc.

The class has a default constructor that initializes data members as follows:

Data Member Type	Default Value
`string`	Empty string
`enum`	Zero
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	Null
`dictionary`	Null
`class/interface`	Null

The constructor won't explicitly initialize a data member if the default C# behavior for that type produces the desired results.

If you wish to ensure that data members of primitive and enumerated types are initialized to specific values, you can declare default values in your Slice definition. The default constructor initializes each of these data members to its declared value instead.

The class also provides a second constructor that has one parameter for each data member. This allows you to construct and initialize a class instance in a single statement:

C#

```
Point p = new Point(5.1, 7.8);
```

The `Clone` method performs a shallow memberwise copy, and the comparison methods have the usual semantics (they perform value comparison).

Note that you can influence the mapping for structures only at the point of definition of a structure, that is, for a particular structure type, you must decide whether you want to use the structure or the class mapping. (You cannot override the structure mapping elsewhere, for example, for individual structure members or operation parameters.)

As we mentioned previously, if a Slice structure (recursively) contains a member of reference type, it is automatically mapped to a C# class. (The compiler behaves as if you had explicitly specified the `"clr:class"` metadata directive for the structure.)

Here is our Employee structure once more:

Slice

```
struct Employee {
    long number;
    string firstName;
    string lastName;
};
```

The structure contains two strings, which are reference types, so the Slice-to-C# compiler generates a C# class for this structure:

Copyright 2016, ZeroC, Inc.

C#

```
public partial class Employee : _System.ICloneable
{
    public long number;
    public string firstName;
    public string lastName;

    public Employee();
    public Employee(long number, string firstName, string lastName);

    public object Clone();

    public override int GetHashCode();
    public override bool Equals(object other);

    public static bool operator==(Employee lhs, Employee rhs);
    public static bool operator!=(Employee lhs, Employee rhs);
}
```

Property Mapping for Structures in C#

You can instruct the compiler to emit property definitions instead of public data members. For example:

Slice

```
["clr:property"] struct Point {
    double x;
    double y;
};
```

The `"clr:property"` metadata directive causes the compiler to generate a property for each Slice data member:

Copyright 2016, ZeroC, Inc.

C#

```
public partial struct Point
{
    private double x_prop;
    public double x {
        get {
            return x_prop;
        }
        set {
            x_prop = value;
        }
    }

    private double y_prop;
    public double y {
        get {
            return y_prop;
        }
        set {
            y_prop = value;
        }
    }

    // Other methods here...
}
```

Note that the properties are non-virtual because C# structures cannot have virtual properties. However, if you apply the `"clr:property"` directive to a structure that contains a member of reference type, or if you combine the `"clr:property"` and `"clr:class"` directives, the generated properties are virtual. For example:

Slice

```
["clr:property", "clr:class"]
struct Point {
    double x;
    double y;
};
```

This generates the following code:

Copyright 2016, ZeroC, Inc.

C#

```
public partial class Point : System.ICloneable
{
    private double x_prop;
    public virtual double x {
        get {
            return x_prop;
        }
        set {
            x_prop = value;
        }
    }

    private double y_prop;
    public virtual double y {
        get {
            return y_prop;
        }
        set {
            y_prop = value;
        }
    }

    // Other methods here...
}
```

See Also

- Metadata
- C-Sharp Mapping for Identifiers
- C-Sharp Mapping for Modules
- C-Sharp Mapping for Built-In Types
- C-Sharp Mapping for Enumerations
- C-Sharp Mapping for Sequences
- C-Sharp Mapping for Dictionaries
- C-Sharp Collection Comparison
- C-Sharp Mapping for Constants
- C-Sharp Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Sequences

Ice for .NET supports several different mappings for sequences. By default, sequences are mapped to arrays. You can use metadata directives to map sequences to a number of alternative types:

- `System.Collections.Generic.List`
- `System.Collections.Generic.LinkedList`
- `System.Collections.Generic.Queue`
- `System.Collections.Generic.Stack`
- User-defined custom types that derive from `System.Collections.Generic.IEnumerable<T>`.

The different mappings allow you to map a sequence to a container type that provides the correct performance trade-off for your application.

On this page:

- Array Mapping for Sequences in C#
- Mapping to Predefined Generic Containers for Sequences in C#
- Mapping to Custom Types for Sequences in C#
- Multi-Dimensional Sequences in C#

Array Mapping for Sequences in C#

By default, the Slice-to-C# compiler maps sequences to arrays. Interestingly, no code is generated in this case; you simply define an array of elements to model the Slice sequence. For example:

Slice

```
sequence<Fruit> FruitPlatter;
```

Given this definition, to create a sequence containing an apple and an orange, you could write:

C#

```
Fruit[] fp = { Fruit.Apple, Fruit.Orange };
```

Or, alternatively:

C#

```
Fruit fp[] = new Fruit[2];
fp[0] = Fruit.Apple;
fp[1] = Fruit.Orange;
```

The array mapping for sequences is both simple and efficient, especially for sequences that do not need to provide insertion or deletion other than at the end of the sequence.

Mapping to Predefined Generic Containers for Sequences in C#

With metadata directives, you can change the default mapping for sequences to use generic containers provided by .NET. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
["clr:generic:List"] sequence<string> StringSeq;
["clr:generic:LinkedList"] sequence<Fruit> FruitSeq;
["clr:generic:Queue"] sequence<int> IntQueue;
["clr:generic:Stack"] sequence<double> DoubleStack;
```

The "`clr:generic:`*type*" metadata directive causes the `slice2cs` compiler to the map the corresponding sequence to one of the containers in the `System.Collections.Generic` namespace. For example, the `Queue` sequence maps to `System.Collections.Generic.Queue<int>` due to its metadata directive.

The predefined containers allow you to select an appropriate space-performance trade-off, depending on how your application uses a sequence. In addition, if a sequence contains value types, such as `int`, the generic containers do not incur the cost of boxing and unboxing and so are quite efficient. (For example, `System.Collections.Generic.List<int>` performs within a few percentage points of an integer array for insertion and deletion at the end of the sequence, but has the advantage of providing a richer set of operations.)

Generic containers can be used for sequences of any element type except objects. For sequences of objects, only `List` is supported because it provides the functionality required for efficient unmarshaling. Metadata that specifies any other generic type is ignored with a warning:

Slice

```
class MyClass {
    // ...
};

["clr:generic:List"]
sequence<MyClass> MyClassList; // OK

["clr:generic:LinkedList"]
sequence<MyClass> MyClassLinkedList; // Ignored
```

In this example, sequence type `MyClassList` maps to the generic container `System.Collections.Generic.List<MyClass>`, but sequence type `MyClassLinkedList` uses the default array mapping.

Mapping to Custom Types for Sequences in C#

If the array mapping and the predefined containers are unsuitable for your application (for example, because you may need a priority queue, which does not come with .NET), you can implement your own custom containers and direct `slice2cs` to map sequences to these custom containers. For example:

Slice

```
["clr:generic:MyTypes.PriorityQueue"] sequence<int> Queue;
```

This metadata directive causes the Slice `Queue` sequence to be mapped to the type `MyTypes.PriorityQueue`. You must specify the fully-qualified name of your custom type following the `clr:generic:` prefix. This is because the generated code prepends a `global::` qualifier to the type name you provide; for the preceding example, the generated code refers to your custom type as `global::MyTypes.Pri`

Copyright 2016, ZeroC, Inc.

`orityQueue<int>`.

Your custom type can have whatever interface you deem appropriate, but it must meet the following requirements:

- The custom type must derive from `System.Collections.Generic.IEnumerable<T>`.
- The custom type must provide a readable `Count` property that returns the number of elements in the collection.
- The custom type must provide an `Add` method that appends an element to the end of the collection.
- If (and only if) the Slice sequence contains elements that are Slice classes, the custom type must provide an indexer that sets the value of an element at a specific index. (Indexes, as usual, start at zero.)

As an example, here is a minimal class (omitting implementation) that meets these criteria:

C#

```
public class PriorityQueue<T> : IEnumerable<T>
{
    public IEnumerator<T> GetEnumerator();

    public int Count
        get;

    public void Add(T elmt);

    public T this[int index] // Needed for class elements only.
        set;

    // Other methods and data members here...
}
```

Multi-Dimensional Sequences in C#

Slice permits you to define sequences of sequences, for example:

Slice

```
enum Fruit { Apple, Orange, Pear };
["clr:generic:List"] sequence<Fruit> FruitPlatter;
["clr:generic:LinkedList"] sequence<FruitPlatter> Cornucopia;
```

If we use these definitions as shown, the type of FruitPlatter in the generated code is:

C#

```
System.Collections.Generic.LinkedList<System.Collections.Generic.List<Fruit>>
```

Here the outer sequence contains elements of type `List<Fruit>`, as you would expect.

Now let us modify the definition to change the mapping of `FruitPlatter` to an array:

Copyright 2016, ZeroC, Inc.

Slice

```
enum Fruit { Apple, Orange, Pear };
sequence<Fruit> FruitPlatter;
["clr:LinkedList"] sequence<FruitPlatter> Cornucopia;
```

With this definition, the type of `Cornucopia` becomes:

C#

```
System.Collections.Generic.LinkedList<Fruit[]>
```

The generated code now no longer mentions the type `FruitPlatter` anywhere and deals with the outer sequence elements as an array of `Fruit` instead.

See Also

- Metadata
- C-Sharp Mapping for Identifiers
- C-Sharp Mapping for Modules
- C-Sharp Mapping for Built-In Types
- C-Sharp Mapping for Enumerations
- C-Sharp Mapping for Structures
- C-Sharp Mapping for Dictionaries
- C-Sharp Collection Comparison
- C-Sharp Mapping for Constants
- C-Sharp Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Dictionaries

Ice for .NET supports two mappings for dictionaries. By default, a dictionary maps to `System.Collections.Generic.Dictionary<T>`. You can use a metadata directive to map a dictionary to System.Collections.Generic.SortedDictionary instead.

On this page:

- Default Mapping for Dictionaries in C#
- Custom Mapping for Dictionaries in C#

Default Mapping for Dictionaries in C#

Here is the definition of our EmployeeMap once more:

Slice

```
dictionary<long, Employee> EmployeeMap;
```

By default, the Slice-to-C# compiler maps the dictionary to the following type:

C#

```
System.Collections.Generic.Dictionary<long, Employee>
```

Custom Mapping for Dictionaries in C#

You can use the `"clr:generic:SortedDictionary"` metadata directive to change the default mapping to use a sorted dictionary instead:

Slice

```
["clr:generic:SortedDictionary"]
dictionary<long, Employee> EmployeeMap;
```

With this definition, the type of the dictionary becomes:

C#

```
System.Collections.Generic.SortedDictionary<long, Employee>
```

See Also

- Metadata
- C-Sharp Mapping for Identifiers
- C-Sharp Mapping for Modules
- C-Sharp Mapping for Built-In Types
- C-Sharp Mapping for Enumerations
- C-Sharp Mapping for Structures
- C-Sharp Mapping for Sequences
- C-Sharp Collection Comparison

Copyright 2016, ZeroC, Inc.

- C-Sharp Mapping for Constants
- C-Sharp Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

C-Sharp Collection Comparison

The utility class `Ice.CollectionComparer` allows you to compare collections for equality:

C#

```
public class CollectionComparer {
    public static bool
    Equals(System.Collections.IDictionary d1, System.Collections.IDictionary d2);

    public static bool
    Equals(System.Collections.ICollection c1, System.Collections.ICollection c2);

    public static bool
    Equals(System.Collections.IEnumerable c1, System.Collections.IEnumerable c2);
}
```

Equality of the elements in a collection is determined by calling the elements' `Equals` method.

Two dictionaries are equal if they contain the same number of entries with identical keys and values.

Two collections that derive from `ICollection` or `IEnumerable` are equal if they contain the same number of entries and entries compare equal. Note that order is significant, so corresponding entries must not only be equal but must also appear in the same position.

See Also

- C-Sharp Mapping for Identifiers
- C-Sharp Mapping for Modules
- C-Sharp Mapping for Built-In Types
- C-Sharp Mapping for Enumerations
- C-Sharp Mapping for Structures
- C-Sharp Mapping for Sequences
- C-Sharp Mapping for Dictionaries
- C-Sharp Mapping for Constants
- C-Sharp Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Constants

Here are the sample constant definitions once more:

Slice

```
const bool      AppendByDefault = true;
const byte      LowerNibble = 0x0f;
const string    Advice = "Don't Panic!";
const short     TheAnswer = 42;
const double    PI = 3.1416;

enum Fruit { Apple, Pear, Orange };
const Fruit     FavoriteFruit = Pear;
```

Here are the generated definitions for these constants:

C#

```
public abstract class AppendByDefault
{
    public const bool value = true;
}

public abstract class LowerNibble
{
    public const byte value = 15;
}

public abstract class Advice
{
    public const string value = "Don't Panic!";
}

public abstract class TheAnswer
{
    public const short value = 42;
}

public abstract class PI
{
    public const double value = 3.1416;
}

public enum Fruit { Apple, Pear, Orange }

public abstract class FavoriteFruit
{
    public const Fruit value = Fruit.Pear;
}
```

As you can see, each Slice constant is mapped to a class with the same name as the constant. The class contains a member named `value` that holds the value of the constant.

Copyright 2016, ZeroC, Inc.

The mapping to classes instead of to plain constants is necessary because C# does not permit constant definitions at namespace scope.

Slice string literals that contain non-ASCII characters or universal character names are mapped to C# string literals with universal character names. For example:

Slice

```
const string Egg = "œuf";
const string Heart = "c\u0153ur";
const string Banana = "\U0001F34C";
```

is mapped to:

C++

```
public abstract class Egg
{
    public const string value = "\u0153uf";
}

public abstract class Heart
{
    public const string value = "c\u0153ur";
}

public abstract class Banana
{
    public const string value = "\ud83c\udf4c";
}
```

See Also

- Constants and Literals
- C-Sharp Mapping for Identifiers
- C-Sharp Mapping for Modules
- C-Sharp Mapping for Built-In Types
- C-Sharp Mapping for Enumerations
- C-Sharp Mapping for Structures
- C-Sharp Mapping for Sequences
- C-Sharp Mapping for Dictionaries
- C-Sharp Collection Comparison
- C-Sharp Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Exceptions

On this page:

- Inheritance Hierarchy for Exceptions in C#
- C# Mapping for User Exceptions
- C# Default Constructors for User Exceptions
- C# Mapping for Run-Time Exceptions

Inheritance Hierarchy for Exceptions in C#

The mapping for exceptions is based on the inheritance hierarchy shown below:

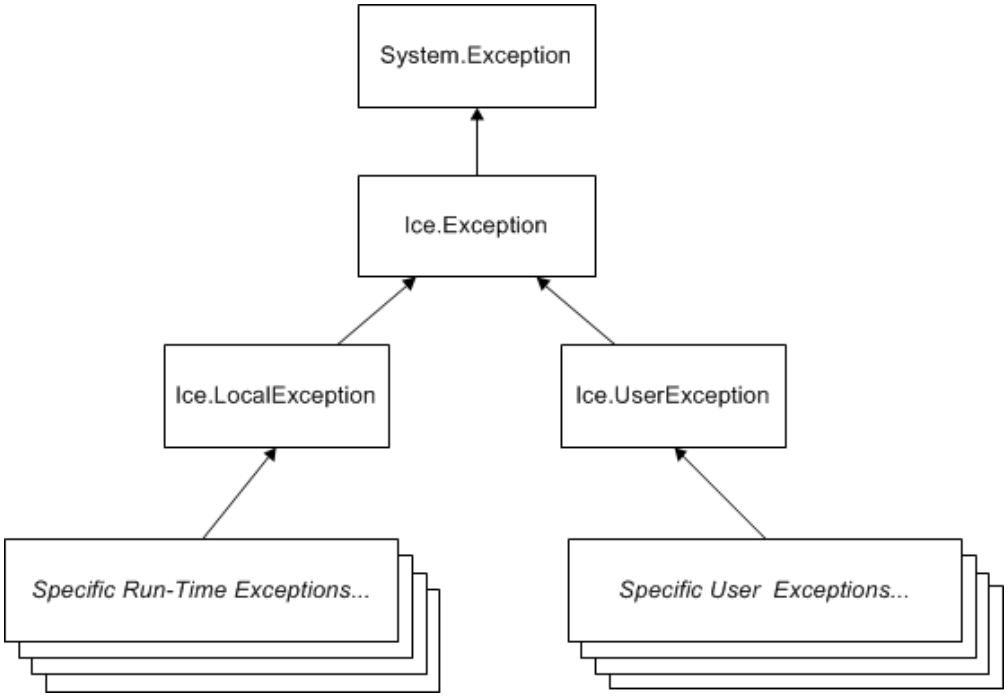

Inheritance structure for exceptions.

The ancestor of all exceptions is `System.Exception`. Derived from that is `Ice.Exception`, which provides the definitions of a number of constructors. `Ice.LocalException` and `Ice.UserException` are derived from `Ice.Exception` and form the base of all run-time and user exceptions, respectively.

The constructors defined in `Ice.Exception` have the following signatures:

C#

```
public abstract class Exception : System.Exception
{
    public Exception();
    public Exception(System.Exception ex);
}
```

Each concrete derived exception class implements these constructors. The second constructor initializes the `InnerException` property of `System.Exception`. (Both constructors set the `Message` property to the empty string.)

C# Mapping for User Exceptions

Copyright 2016, ZeroC, Inc.

Here is a fragment of the Slice definition for our world time server once more:

Slice

```
exception GenericError {
    string reason;
};

exception BadTimeVal extends GenericError {};

exception BadZoneName extends GenericError {};
```

These exception definitions map as follows:

C#

```
public partial class GenericError : Ice.UserException
{
    public string reason;

    public GenericError();
    public GenericError(System.Exception ex__);
    public GenericError(string reason);
    public GenericError(string reason, System.Exception ex__);

    // GetHashCode and comparison methods defined here,
    // as well as mapping-internal methods.
}

public partial class BadTimeVal : M.GenericError
{
    public BadTimeVal();
    public BadTimeVal(System.Exception ex__);
    public BadTimeVal(string reason);
    public BadTimeVal(string reason, System.Exception ex__);

    // GetHashCode and comparison methods defined here,
    // as well as mapping-internal methods.
}

public partial class BadZoneName : M.GenericError
{
    public BadZoneName();
    public BadZoneName(System.Exception ex__);
    public BadZoneName(string reason);
    public BadZoneName(string reason, System.Exception ex__);

    // GetHashCode and comparison methods defined here,
    // as well as mapping-internal methods.
}
```


Copyright 2016, ZeroC, Inc.

Each Slice exception is mapped to a C# partial class with the same name. For each exception member, the corresponding class contains a public data member. (Obviously, because `BadTimeVal` and `BadZoneName` do not have members, the generated classes for these exceptions also do not have members.) Optional data members are mapped to instances of the `Ice.Optional` type.

The inheritance structure of the Slice exceptions is preserved for the generated classes, so `BadTimeVal` and `BadZoneName` inherit from `GenericError`.

All user exceptions are derived from the base class `Ice.UserException`. This allows you to catch all user exceptions generically by installing a handler for `Ice.UserException`. Similarly, you can catch all Ice run-time exceptions with a handler for `Ice.LocalException`, and you can catch all Ice exceptions with a handler for `Ice.Exception`.

All exceptions provide the usual `GetHashCode` and `Equals` methods, as well as the `==` and `!=` comparison operators.

The generated exception classes also contain other member functions that are not shown here; these member functions are internal to the C# mapping and are not meant to be called by application code.

C# Default Constructors for User Exceptions

An exception has a default constructor that initializes data members as follows:

Data Member Type	Default Value
`string`	Empty string
`enum`	Zero
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	Null
`dictionary`	Null
`class/interface`	Null

The constructor won't explicitly initialize a data member if the default C# behavior for that type produces the desired results.

If you wish to ensure that data members of primitive and enumerated types are initialized to specific values, you can declare default values in your Slice definition. The default constructor initializes each of these data members to its declared value instead.

An exception also provides a constructor that accepts one parameter for each data member so that you can construct and initialize a class instance in a single statement (instead of first having to construct the instance and then assign to its members). For a derived exception, this constructor accepts one argument for each base exception member, plus one argument for each derived exception member, in base-to-derived order. For each optional data member, the constructor accepts an `Ice.Optional` parameter of the appropriate type.

C# Mapping for Run-Time Exceptions

The Ice run time throws run-time exceptions for a number of pre-defined error conditions. All run-time exceptions directly or indirectly derive from `Ice.LocalException` (which, in turn, derives from `Ice.Exception`).

`Ice.LocalException` implements a `Clone` method that is inherited by its derived exceptions, so you can make memberwise shallow copies of exceptions.

By catching exceptions at the appropriate point in the inheritance hierarchy, you can handle exceptions according to the category of error they indicate:

- `Ice.Exception`
  This is the root of the inheritance tree for both run-time and user exceptions.
- `Ice.LocalException`
  This is the root of the inheritance tree for run-time exceptions.
- `Ice.UserException`
  This is the root of the inheritance tree for user exceptions.

Copyright 2016, ZeroC, Inc.

- `Ice.TimeoutException`
  This is the base exception for both operation-invocation and connection-establishment timeouts.
- `Ice.ConnectTimeoutException`
  This exception is raised when the initial attempt to establish a connection to a server times out.

For example, a `ConnectTimeoutException` can be handled as `ConnectTimeoutException`, `TimeoutException`, `LocalException`, or `Exception`.

You will probably have little need to catch run-time exceptions as their most-derived type and instead catch them as `LocalException`; the fine-grained error handling offered by the remainder of the hierarchy is of interest mainly in the implementation of the Ice run time. Exceptions to this rule are the exceptions related to facet and object life cycles, which you may want to catch explicitly. These exceptions are `FacetNotExistException` and `ObjectNotExistException`, respectively.

See Also

- User Exceptions
- Run-Time Exceptions
- C-Sharp Mapping for Identifiers
- C-Sharp Mapping for Modules
- C-Sharp Mapping for Built-In Types
- C-Sharp Mapping for Enumerations
- C-Sharp Mapping for Structures
- C-Sharp Mapping for Sequences
- C-Sharp Mapping for Dictionaries
- C-Sharp Collection Comparison
- C-Sharp Mapping for Constants
- C-Sharp Mapping for Optional Values
- Versioning
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Interfaces

The mapping of Slice interfaces revolves around the idea that, to invoke a remote operation, you call a member function on a local class instance that is a proxy for the remote object. This makes the mapping easy and intuitive to use because making a remote procedure call is no different from making a local procedure call (apart from error semantics).

On this page:

- Proxy Interfaces in C#
- The Ice.ObjectPrx Interface in C#
- Proxy Helpers in C#
- Using Proxy Methods in C#
- Object Identity and Proxy Comparison in C#

Proxy Interfaces in C#

On the client side, a Slice interface maps to a C# interface with member functions that correspond to the operations on that interface. Consider the following simple interface:

Slice

```
interface Simple {
    void op();
};
```

The Slice compiler generates the following definition for use by the client:

C#

```
public interface SimplePrx : Ice.ObjectPrx
{
    void op();
    void op(System.Collections.Generic.Dictionary<string, string> __context);
}
```

As you can see, the compiler generates a *proxy interface* `SimplePrx`. In general, the generated name is `<`*`interface-name`*`>Prx`. If an interface is nested in a module `M`, the generated interface is part of namespace `M`, so the fully-qualified name is `M.<`*`interface-name`*`>Prx`.

In the client's address space, an instance of `SimplePrx` is the local ambassador for a remote instance of the `Simple` interface in a server and is known as a *proxy instance*. All the details about the server-side object, such as its address, what protocol to use, and its object identity are encapsulated in that instance.

Note that `SimplePrx` inherits from `Ice.ObjectPrx`. This reflects the fact that all Ice interfaces implicitly inherit from `Ice::Object`.

For each operation in the interface, the proxy class has a member function of the same name. For the preceding example, we find that the operation `op` has been mapped to the method `op`. Also note that `op` is overloaded: the second version of `op` has a parameter `__context`, which is a dictionary of string pairs. This parameter is for use by the Ice run time to store information about how to deliver a request. You normally do not need to use it. (We examine the `__context` parameter in detail in Request Contexts. The parameter is also used by IceStorm.)

Because all the `<`*`interface-name`*`>Prx` types are interfaces, you cannot instantiate an object of such a type. Instead, proxy instances are always instantiated on behalf of the client by the Ice run time, so client code never has any need to instantiate a proxy directly.The proxy references handed out by the Ice run time are always of type `<`*`interface-name`*`>Prx`; the concrete implementation of the interface is part of the Ice run time and does not concern application code.

A value of `null` denotes the null proxy. The null proxy is a dedicated value that indicates that a proxy points "nowhere" (denotes no object).

Copyright 2016, ZeroC, Inc.

The `Ice.ObjectPrx` Interface in C#

All Ice objects have `Object` as the ultimate ancestor type, so all proxies inherit from `Ice.ObjectPrx`. `ObjectPrx` provides a number of methods:

C#

```
namespace Ice
{
    public interface ObjectPrx
    {
        Identity ice_getIdentity();
        bool ice_isA(string id);
        string ice_id();
        void ice_ping();

        int GetHashCode();
        bool Equals(object r);

        // Defined in a helper class:
        //
        public static bool Equals(Ice.ObjectPrx lhs, ObjectPrx rhs);
        public static bool operator==(ObjectPrx lhs, ObjectPrx rhs);
        public static bool operator!=(ObjectPrx lhs, ObjectPrx rhs);

        // ...
    }
}
```

Note that the static methods are not actually defined in `Ice.ObjectPrx`, but in a helper class that becomes a base class of an instantiated proxy. However, this is simply an internal detail of the C# mapping — conceptually, these methods belong with `Ice.ObjectPrx`, so we discuss them here.

The methods behave as follows:

- **`ice_getIdentity`**
  This method returns the identity of the object denoted by the proxy. The identity of an Ice object has the following Slice type:

  Slice

  ```
  module Ice {
      struct Identity {
          string name;
          string category;
      };
  };
  ```

  To see whether two proxies denote the same object, first obtain the identity for each object and then compare the identities:

Copyright 2016, ZeroC, Inc.

C#

```
Ice.ObjectPrx o1 = ...;
Ice.ObjectPrx o2 = ...;
Ice.Identity i1 = o1.ice_getIdentity();
Ice.Identity i2 = o2.ice_getIdentity();

if (i1.Equals(i2))
    // o1 and o2 denote the same object
else
    // o1 and o2 denote different objects
```

- **ice_isA**
  The ice_isA method determines whether the object denoted by the proxy supports a specific interface. The argument to ice_isA is a type ID. For example, to see whether a proxy of type ObjectPrx denotes a Printer object, we can write:

  C#

  ```
  Ice.ObjectPrx o = ...;
  if (o != null && o.ice_isA("::Printer"))
      // o denotes a Printer object
  else
      // o denotes some other type of object
  ```

  Note that we are testing whether the proxy is null before attempting to invoke the ice_isA method. This avoids getting a NullReferenceException if the proxy is null.

- **ice_ids**
  The ice_ids method returns an array of strings representing all of the type IDs that the object denoted by the proxy supports.

- **ice_id**
  The ice_id method returns the type ID of the object denoted by the proxy. Note that the type returned is the type of the actual object, which may be more derived than the static type of the proxy. For example, if we have a proxy of type BasePrx, with a static type ID of ::Base, the return value of ice_id might be ::Base, or it might something more derived, such as ::Derived.

- **ice_ping**
  The ice_ping method provides a basic reachability test for the object. If the object can physically be contacted (that is, the object exists and its server is running and reachable), the call completes normally; otherwise, it throws an exception that indicates why the object could not be reached, such as ObjectNotExistException or ConnectTimeoutException.

- **Equals**
  This method compares two proxies for equality. Note that all aspects of proxies are compared by this operation, such as the communication endpoints for the proxy. This means that, in general, if two proxies compare unequal, that does *not* imply that they denote different objects. For example, if two proxies denote the same Ice object via different transport endpoints, equals returns false even though the proxies denote the same object.

The ice_isA, ice_ids, ice_id, and ice_ping methods are remote operations and therefore support an additional overloading that accepts a request context. Also note that there are other methods in ObjectPrx, not shown here. These methods provide different ways to dispatch a call and also provide access to an object's facets.

Proxy Helpers in C#

For each Slice interface, apart from the proxy interface, the Slice-to-C# compiler creates a helper class: for an interface Simple, the name of

Copyright 2016, ZeroC, Inc.

the generated helper class is `SimplePrxHelper`.

You can ignore the `ObjectPrxHelperBase` base class — it exists for mapping-internal purposes.

The helper class contains two methods of interest:

C#

```
public class SimplePrxHelper : Ice.ObjectPrxHelperBase, SimplePrx
{
    public static SimplePrx checkedCast(Ice.ObjectPrx b);
    public static SimplePrx checkedCast(
        Ice.ObjectPrx b,
        System.Collections.Generic.Dictionary<string, string> ctx);
    public static SimplePrx uncheckedCast(Ice.ObjectPrx b);
    public static string ice_staticId();

    // ...
}
```

For `checkedCast`, if the passed proxy is for an object of type `Simple`, or a proxy for an object with a type derived from `Simple`, the cast returns a non-null reference to a proxy of type `SimplePrx`; otherwise, if the passed proxy denotes an object of a different type (or if the passed proxy is null), the cast returns a null reference.

Given a proxy of any type, you can use a `checkedCast` to determine whether the corresponding object supports a given type, for example:

C#

```
Ice.ObjectPrx obj = ...;        // Get a proxy from somewhere...

SimplePrx simple = SimplePrxHelper.checkedCast(obj);
if (simple != null)
    // Object supports the Simple interface...
else
    // Object is not of type Simple...
```

Note that a `checkedCast` contacts the server. This is necessary because only the implementation of an object in the server has definite knowledge of the type of an object. As a result, a `checkedCast` may throw a `ConnectTimeoutException` or an `ObjectNotExistException`. (This also explains the need for the helper class: the Ice run time must contact the server, so we cannot use a C# down-cast.)

In contrast, an `uncheckedCast` does not contact the server and unconditionally returns a proxy of the requested type. However, if you do use an `uncheckedCast`, you must be certain that the proxy really does support the type you are casting to; otherwise, if you get it wrong, you will most likely get a run-time exception when you invoke an operation on the proxy. The most likely error for such a type mismatch is `OperationNotExistException`. However, other exceptions, such as a marshaling exception are possible as well. And, if the object happens to have an operation with the correct name, but different parameter types, no exception may be reported at all and you simply end up sending the invocation to an object of the wrong type; that object may do rather nonsensical things. To illustrate this, consider the following two interfaces:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Process {
    void launch(int stackSize, int dataSize);
};

// ...

interface Rocket {
    void launch(float xCoord, float yCoord);
};
```

Suppose you expect to receive a proxy for a `Process` object and use an `uncheckedCast` to down-cast the proxy:

C#

```
Ice.ObjectPrx obj = ...;                  // Get proxy...
ProcessPrx process = ProcessPrxHelper.uncheckedCast(obj); // No worries...
process.launch(40, 60);                   // Oops...
```

If the proxy you received actually denotes a `Rocket` object, the error will go undetected by the Ice run time: because `int` and `float` have the same size and because the Ice protocol does not tag data with its type on the wire, the implementation of `Rocket::launch` will simply misinterpret the passed integers as floating-point numbers.

In fairness, this example is somewhat contrived. For such a mistake to go unnoticed at run time, both objects must have an operation with the same name and, in addition, the run-time arguments passed to the operation must have a total marshaled size that matches the number of bytes that are expected by the unmarshaling code on the server side. In practice, this is extremely rare and an incorrect `uncheckedCast` typically results in a run-time exception.

A final warning about down-casts: you must use either a `checkedCast` or an `uncheckedCast` to down-cast a proxy. If you use a C# cast, the behavior is undefined.

Another method defined by every helper class is `ice_staticId`, which returns the type ID string corresponding to the interface. As an example, for the Slice interface `Simple` in module `M`, the string returned by `ice_staticId` is `"::M::Simple"`.

Using Proxy Methods in C#

The base proxy class `ObjectPrx` supports a variety of methods for customizing a proxy. Since proxies are immutable, each of these "factory methods" returns a copy of the original proxy that contains the desired modification. For example, you can obtain a proxy configured with a ten second timeout as shown below:

C#

```
Ice.ObjectPrx proxy = communicator.stringToProxy(...);
proxy = proxy.ice_timeout(10000);
```

A factory method returns a new proxy object if the requested modification differs from the current proxy, otherwise it returns the current proxy. With few exceptions, factory methods return a proxy of the same type as the current proxy, therefore it is generally not necessary to repeat a `checkedCast` or `uncheckedCast` after using a factory method. However, a regular cast is still required, as shown in the example below:

Copyright 2016, ZeroC, Inc.

C#

```
Ice.ObjectPrx base = communicator.stringToProxy(...);
HelloPrx hello = HelloPrxHelper.checkedCast(base);
hello = (HelloPrx)hello.ice_timeout(10000); # Type is preserved
hello.sayHello();
```

The only exceptions are the factory methods `ice_facet` and `ice_identity`. Calls to either of these methods may produce a proxy for an object of an unrelated type, therefore they return a base proxy that you must subsequently down-cast to an appropriate type.

Object Identity and Proxy Comparison in C#

Proxies provide an `Equals` method that compares proxies:

C#

```
public interface ObjectPrx {
    bool Equals(object r);
}
```

Note that proxy comparison with `Equals` uses *all* of the information in a proxy for the comparison. This means that not only the object identity must match for a comparison to succeed, but other details inside the proxy, such as the protocol and endpoint information, must be the same. In other words, comparison with `Equals` (or == and `!=`) tests for *proxy* identity, *not* object identity. A common mistake is to write code along the following lines:

C#

```
Ice.ObjectPrx p1 = ...;        // Get a proxy...
Ice.ObjectPrx p2 = ...;        // Get another proxy...

if (!p1.Equals(p2)) {
    // p1 and p2 denote different objects       // WRONG!
} else {
    // p1 and p2 denote the same object         // Correct
}
```

Even though `p1` and `p2` differ, they may denote the same Ice object. This can happen because, for example, both `p1` and `p2` embed the same object identity, but each use a different protocol to contact the target object. Similarly, the protocols may be the same, but denote different endpoints (because a single Ice object can be contacted via several different transport endpoints). In other words, if two proxies compare equal with `Equals`, we know that the two proxies denote the same object (because they are identical in all respects); however, if two proxies compare unequal with `Equals`, we know absolutely nothing: the proxies may or may not denote the same object.

To compare the object identities of two proxies, you can use a helper function in the `Ice.Util` class:

Copyright 2016, ZeroC, Inc.

C#

```
public sealed class Util {
    public static int proxyIdentityCompare(ObjectPrx lhs, ObjectPrx rhs);
    public static int proxyIdentityAndFacetCompare(ObjectPrx lhs, ObjectPrx rhs);
// ...
```

`proxyIdentityCompare` allows you to correctly compare proxies for identity:

C#

```
Ice.ObjectPrx p1 = ...;        // Get a proxy...
Ice.ObjectPrx p2 = ...;        // Get another proxy...

if (Ice.Util.proxyIdentityCompare(p1, p2) != 0) {
    // p1 and p2 denote different objects       // Correct
} else {
    // p1 and p2 denote the same object         // Correct
}
```

The function returns 0 if the identities are equal, `-1` if `p1` is less than `p2`, and 1 if `p1` is greater than `p2`. (The comparison uses `name` as the major and `category` as the minor sort key.)

The `proxyIdentityAndFacetCompare` function behaves similarly, but compares both the identity and the facet name.

The C# mapping also provides two helper classes in the `Ice` namespace that allow you to insert proxies into hashtables or ordered collections, based on the identity, or the identity plus the facet name:

C#

```
public class ProxyIdentityKey
    : System.Collections.IHashCodeProvider,
      System.Collections.IComparer {

    public int GetHashCode(object obj);
    public int Compare(object obj1, object obj2);
}

public class ProxyIdentityFacetKey
    : System.Collections.IHashCodeProvider,
      System.Collections.IComparer {

    public int GetHashCode(object obj);
    public int Compare(object obj1, object obj2);
}
```

Note these classes derive from `IHashCodeProvider` and `IComparer`, so they can be used for both hash tables and ordered collections.

Copyright 2016, ZeroC, Inc.

See Also

- Interfaces, Operations, and Exceptions
- Proxies for Ice Objects
- C-Sharp Mapping for Operations
- Operations on Object
- Proxy Methods
- Versioning
- IceStorm

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Operations

On this page:

- Basic C# Mapping for Operations
- Normal and idempotent Operations in C#
- Passing Parameters in C#
  - In-Parameters in C#
  - Out-Parameters in C#
  - Null Parameters in C#
  - Optional Parameters in C#
- Exception Handling in C#
  - Exceptions and Out-Parameters in C#

Basic C# Mapping for Operations

As we saw in the C# mapping for interfaces, for each operation on an interface, the proxy class contains a corresponding member function with the same name. To invoke an operation, you call it via the proxy. For example, here is part of the definitions for our file system:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };
    // ...
};
```

The `name` operation returns a value of type `string`. Given a proxy to an object of type `Node`, the client can invoke the operation as follows:

C#

```
NodePrx node = ...;             // Initialize proxy
string name = node.name();      // Get name via RPC
```

This illustrates the typical pattern for receiving return values: return values are returned by reference for complex types, and by value for simple types (such as `int` or `double`).

Normal and `idempotent` Operations in C#

You can add an `idempotent` qualifier to a Slice operation. As far as the signature for the corresponding proxy method is concerned, `idempotent` has no effect. For example, consider the following interface:

Slice

```
interface Example {
                string op1();
    idempotent  string op2();
};
```

The proxy interface for this is:

Copyright 2016, ZeroC, Inc.

C#

```
public interface ExamplePrx : Ice.ObjectPrx
{
    string op1();
    string op2();
}
```

Because `idempotent` affects an aspect of call dispatch, not interface, it makes sense for the two methods to be mapped the same.

Passing Parameters in C#

In-Parameters in C#

The parameter passing rules for the C# mapping are very simple: parameters are passed either by value (for value types) or by reference (for reference types). Semantically, the two ways of passing parameters are identical: it is guaranteed that the value of a parameter will not be changed by the invocation.

Here is an interface with operations that pass parameters of various types from client to server:

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ClientToServer {
    void op1(int i, float f, bool b, string s);
    void op2(NumberAndString ns, StringSeq ss, StringTable st);
    void op3(ClientToServer* proxy);
};
```

The Slice compiler generates the following proxy for these definitions:

Copyright 2016, ZeroC, Inc.

C#

```
public interface ClientToServerPrx : Ice.ObjectPrx
{
    void op1(int i, float f, bool b, string s);
    void op2(NumberAndString ns, string[] ss, Dictionary<long, string[]> st);
    void op3(ClientToServerPrx proxy);
}
```

Given a proxy to a `ClientToServer` interface, the client code can pass parameters as in the following example:

C#

```
ClientToServerPrx p = ...;              // Get proxy...

p.op1(42, 3.14f, true, "Hello world!"); // Pass simple literals

int i = 42;
float f = 3.14f;
bool b = true;
string s = "Hello world!";
p.op1(i, f, b, s);                      // Pass simple variables

NumberAndString ns = new NumberAndString();
ns.x = 42;
ns.str = "The Answer";
string[] ss = new string[1];
ss[0] = "Hello world!";
Dictionary<long, string[]> st = new Dictionary<long, string[]>();
st[0] = ss;
p.op2(ns, ss, st);                      // Pass complex variables

p.op3(p);                               // Pass proxy
```

Out-Parameters in C#

Slice `out` parameters simply map to C# `out` parameters.

Here again are the same Slice definitions we saw earlier, but this time with all parameters being passed in the `out` direction:

Copyright 2016, ZeroC, Inc.

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ServerToClient {
    void op1(out int i, out float f, out bool b, out string s);
    void op2(out NumberAndString ns, out StringSeq ss, out StringTable st);
    void op3(out ServerToClient* proxy);
};
```

The Slice compiler generates the following code for these definitions:

C#

```
public interface ServerToClientPrx : Ice.ObjectPrx
{
    void op1(out int i, out float f, out bool b, out string s);
    void op2(out NumberAndString ns,
             out string[] ss,
             out Dictionary<long, string[]> st);
    void op3(out ServerToClientPrx proxy);
}
```

Given a proxy to a `ServerToClient` interface, the client code can pass parameters as in the following example:

Copyright 2016, ZeroC, Inc.

C#

```
ClientToServerPrx p = ...;              // Get proxy...

int i;
float f;
bool b;
string s;
p.op1(out i, out f, out b, out s);

NumberAndString ns;
string[] ss;
Dictionary<long, string[]> st;
p.op2(out ns, out ss, out st);

ServerToClientPrx stc;
p.op3(out stc);

System.Console.WriteLine(i);   // Show one of the values
```

Null Parameters in C#

Some Slice types naturally have "empty" or "not there" semantics. Specifically, C# sequences (if mapped to `CollectionBase`), dictionaries, strings, and structures (if mapped to classes) all can be `null`, but the corresponding Slice types do not have the concept of a null value.

- Slice sequences, dictionaries, and strings cannot be null, but can be empty. To make life with these types easier, whenever you pass a C# `null` reference as a parameter or return value of type sequence, dictionary, or string, the Ice run time automatically sends an empty sequence, dictionary, or string to the receiver.
- If you pass a C# `null` reference to a Slice structure that is mapped to a C# class as a parameter or return value, the Ice run time automatically sends a structure whose elements are default-initialized. This means that all proxy members are initialized to `null`, sequence and dictionary members are initialized to empty collections, strings are initialized to the empty string, and members that have a value type are initialized to their default values.

This behavior is useful as a convenience feature: especially for deeply-nested data types, members that are structures, sequences, dictionaries, or strings automatically arrive as an empty value at the receiving end. This saves you having to explicitly initialize, for example, every string element in a large sequence before sending the sequence in order to avoid `NullReferenceException`. Note that using null parameters in this way does *not* create null semantics for Slice sequences, dictionaries, or strings. As far as the object model is concerned, these do not exist (only *empty* sequences, dictionaries, and strings do). For example, whether you send a string as `null` or as an empty string makes no difference to the receiver: either way, the receiver sees an empty string.

Optional Parameters in C#

The mapping for optional parameters is the same as for required parameters, except each optional parameter is encapsulated in an `Ice.Optional` value. Consider the following operation:

Slice

```
optional(1) int execute(optional(2) string params, out optional(3) float value);
```

The C# mapping for this operation is shown below:

Copyright 2016, ZeroC, Inc.

C#

```
Ice.Optional<int> execute(Ice.Optional<string> params, out Ice.Optional<float>
value, ...);
```

The constructors and conversion operators provided by the `Ice.Optional` type simplify the use of optional parameters:

C#

```
Ice.Optional<int> i;
Ice.Optional<float> v;

i = proxy.execute("--file log.txt", out v); // string converted to Optional<string>
i = proxy.execute(Ice.Util.None, out v);    // params is unset

if(v.HasValue)
    Console.WriteLine("value = " + v.Value);
```

A well-behaved program must not assume that an optional parameter always has a value.

Exception Handling in C#

Any operation invocation may throw a run-time exception and, if the operation has an exception specification, may also throw user exceptions. Suppose we have the following simple interface:

Slice

```
exception Tantrum {
    string reason;
};

interface Child {
    void askToCleanUp() throws Tantrum;
};
```

Slice exceptions are thrown as C# exceptions, so you can simply enclose one or more operation invocations in a `try-catch` block:

Copyright 2016, ZeroC, Inc.

C#

```
ChildPrx child = ...;   // Get child proxy...

try
{
    child.askToCleanUp();
}
catch (Tantrum t)
{
    System.Console.Write("The child says: ");
    System.Console.WriteLine(t.reason);
}
```

Typically, you will catch only a few exceptions of specific interest around an operation invocation; other exceptions, such as unexpected run-time errors, will typically be handled by exception handlers higher in the hierarchy. For example:

Copyright 2016, ZeroC, Inc.

C#

```
public class Client
{
    private static void run() {
        ChildPrx child = ...;       // Get child proxy...
        try
        {
            child.askToCleanUp();
        }
        catch (Tantrum t)
        {
            System.Console.Write("The child says: ");
            System.Console.WriteLine(t.reason);
            child.scold();          // Recover from error...
        }
        child.praise();             // Give positive feedback...
    }

    static void Main(string[] args)
    {
        try
        {
            // ...
            run();
            // ...
        }
        catch (Ice.Exception e)
        {
            System.Console.WriteLine(e);
        }
    }
}
```

This code handles a specific exception of local interest at the point of call and deals with other exceptions generically. (This is also the strategy we used for our first simple application.)

Note that the `ToString` method of exceptions prints the name of the exception, any inner exceptions, and the stack trace. Of course, you can be more selective in the way exceptions are displayed. For example, `e.GetType().Name` returns the (unscoped) name of an exception.

Exceptions and Out-Parameters in C#

The Ice run time makes no guarantees about the state of out-parameters when an operation throws an exception: the parameter may still have its original value or may have been changed by the operation's implementation in the target object. In other words, for out-parameters, Ice provides the weak exception guarantee [1] but does not provide the strong exception guarantee.

This is done for reasons of efficiency: providing the strong exception guarantee would require more overhead than can be justified.

See Also

- Operations
- C-Sharp Mapping for Exceptions
- C-Sharp Mapping for Interfaces
- Collocated Invocation and Dispatch

Copyright 2016, ZeroC, Inc.

References

1. Sutter, H. 1999. Exceptional C++: 47 Engineering Puzzles, Programming Problems, and Solutions. Reading, MA: Addison-Wesley.

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Classes

On this page:

- Basic C# Mapping for Classes
- Operations Interfaces in C#
- Inheritance from Ice.Object in C#
- Class Data Members in C#
- Class Operations in C#
- Class Factories in C#
- Class Constructors in C#

Basic C# Mapping for Classes

A Slice class is mapped to a C# class with the same name. By default, the generated class contains a public data member for each Slice data member (just as for structures and exceptions), and a member function for each operation. Alternatively, you can use the property mapping by specifying the `"clr:property"` metadata directive, which generates classes with virtual properties instead of data members.

Consider the following class definition:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Slice compiler generates the following code for this definition:

Copyright 2016, ZeroC, Inc.

C#

```
public interface TimeOfDayOperations_
{
    string format(Ice.Current __current);
}

public interface TimeOfDayOperationsNC_
{
    string format();
}

public abstract partial class TimeOfDay
    : Ice.ObjectImpl,
      TimeOfDayOperations_,
      TimeOfDayOperationsNC_
{
    public short hour;
    public short minute;
    public short second;

    public TimeOfDay()
    {
    }

    public TimeOfDay(short hour, short minute, short second)
    {
        this.hour = hour;
        this.minute = minute;
        this.second = second;
    }

    public string format()
    {
        return format(new Ice.Current());
    }

    public abstract string format(Ice.Current __current);
}
```

There are a number of things to note about the generated code:

1. The compiler generates "operations interfaces" called `TimeOfDayOperations_` and `TimeOfDayOperationsNC_`. These interfaces contain a method for each Slice operation of the class.
2. The generated class `TimeOfDay` inherits (indirectly) from `Ice.Object`. This means that all classes implicitly inherit from `Ice.Object`, which is the ultimate ancestor of all classes. Note that `Ice.Object` is *not* the same as `Ice.ObjectPrx`. In other words, you *cannot* pass a class where a proxy is expected and vice versa.
   If a class has only data members, but no operations, the compiler generates a non-abstract class.
3. The generated class contains a public member for each Slice data member.
4. The generated class inherits member functions for each Slice operation from the operations interfaces.
5. The generated class contains two constructors.

There is quite a bit to discuss here, so we will look at each item in turn.

Copyright 2016, ZeroC, Inc.

Operations Interfaces in C#

The methods in the *<interface-name>*`Operations_` interface have an additional trailing parameter of type `Ice.Current`, whereas the methods in the *<interface-name>*`OperationsNC_` interface lack this additional trailing parameter. The methods without the `Current` parameter simply forward to the methods with a `Current` parameter, supplying a default `Current`. For now, you can ignore this parameter and pretend it does not exist.

If a class has only data members, but no operations, the compiler omits generating the *<interface-name>*`Operations_` and *<interface-name>*`OperationsNC_` interfaces.

Inheritance from `Ice.Object` in C#

Like interfaces, classes implicitly inherit from a common base class, `Ice.Object`. However, as shown in the illustration below, classes inherit from `Ice.Object` instead of `Ice.ObjectPrx` (which is at the base of the inheritance hierarchy for proxies). As a result, you cannot pass a class where a proxy is expected (and vice versa) because the base types for classes and proxies are not compatible.

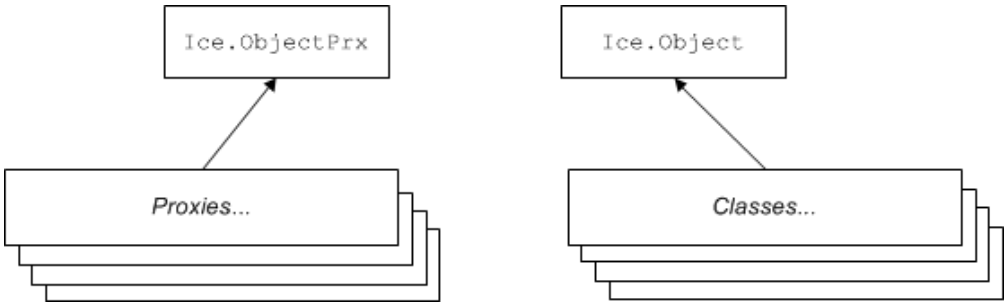

Inheritance from `Ice.ObjectPrx` and `Ice.Object`.

`Ice.Object` contains a number of member functions:

Copyright 2016, ZeroC, Inc.

C#

```
namespace Ice
{
    public interface Object : System.ICloneable
    {
        bool ice_isA(string s);
        bool ice_isA(string s, Current current);

        void ice_ping();
        void ice_ping(Current current);

        string[] ice_ids();
        string[] ice_ids(Current current);

        string ice_id();
        string ice_id(Current current);

        void ice_preMarshal();
        void ice_postUnmarshal();

        DispatchStatus ice_dispatch(Request request,
DispatchInterceptorAsyncCallback cb);
    }
}
```

The member functions of `Ice.Object` behave as follows:

- `ice_isA`
  This function returns `true` if the object supports the given type ID, and `false` otherwise.
- `ice_ping`
  As for interfaces, `ice_ping` provides a basic reachability test for the class.
- `ice_ids`
  This function returns a string sequence representing all of the type IDs supported by this object, including `::Ice::Object`.
- `ice_id`
  This function returns the actual run-time type ID for a class. If you call `ice_id` through a reference to a base instance, the returned type id is the actual (possibly more derived) type ID of the instance.
- `ice_preMarshal`
  The Ice run time invokes this function prior to marshaling the object's state, providing the opportunity for a subclass to validate its declared data members.
- `ice_postUnmarshal`
  The Ice run time invokes this function after unmarshaling an object's state. A subclass typically overrides this function when it needs to perform additional initialization using the values of its declared data members.
- `ice_dispatch`
  This function dispatches an incoming request to a servant. It is used in the implementation of dispatch interceptors.

Note that the generated class does *not* override `GetHashCode` and `Equals`. This means that classes are compared using shallow reference equality, not value equality (as is used for structures).

The class also provides a `Clone` method (whose implementation is inherited from `Ice.ObjectImpl`); the `Clone` method returns a shallow memberwise copy.

Copyright 2016, ZeroC, Inc.

Class Data Members in C#

By default, data members of classes are mapped exactly as for structures and exceptions: for each data member in the Slice definition, the generated class contains a corresponding public data member. Optional data members are mapped to instances of the `Ice.Optional` type .

If you wish to restrict access to a data member, you can modify its visibility using the `protected` metadata directive. The presence of this directive causes the Slice compiler to generate the data member with protected visibility. As a result, the member can be accessed only by the class itself or by one of its subclasses. For example, the `TimeOfDay` class shown below has the `protected` metadata directive applied to each of its data members:

Slice

```
class TimeOfDay {
    ["protected"] short hour;   // 0 - 23
    ["protected"] short minute; // 0 - 59
    ["protected"] short second; // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Slice compiler produces the following generated code for this definition:

C#

```
public abstract partial class TimeOfDay
    : Ice.ObjectImpl,
      TimeOfDayOperations_,
      TimeOfDayOperationsNC_
{
    protected short hour;
    protected short minute;
    protected short second;

    public TimeOfDay()
    {
    }

    public TimeOfDay(short hour, short minute, short second)
    {
        this.hour = hour;
        this.minute = minute;
        this.second = second;
    }

    // ...
}
```

For a class in which all of the data members are protected, the metadata directive can be applied to the class itself rather than to each member individually. For example, we can rewrite the `TimeOfDay` class as follows:

Copyright 2016, ZeroC, Inc.

Slice

```
["protected"] class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

If a protected data member also has the `clr:property` directive, the generated property has protected visibility. Consider the `TimeOfDay` class once again:

Slice

```
["protected", "clr:property"] class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The effects of combining these two metadata directives are shown in the generated code below:

C#

```
public abstract partial class TimeOfDay
    : Ice.ObjectImpl,
      TimeOfDayOperations_,
      TimeOfDayOperationsNC_
{
    private short hour_prop;
    protected short hour {
        get {
            return hour_prop;
        }
        set {
            hour_prop = value;
        }
    }

    // ...
}
```

Refer to the structure mapping for more information on the property mapping for data members.

Class Operations in C#

Copyright 2016, ZeroC, Inc.

Operations of classes are mapped to abstract member functions in the generated class. This means that, if a class contains operations (such as the `format` operation of our `TimeOfDay` class), you must provide an implementation of the operation in a class that is derived from the generated class. For example:

C#

```
public class TimeOfDayI : TimeOfDay
{
    public string format(Ice.Current current)
    {
        return   hour.ToString("D2") + ":"
               + minute.ToString("D2") + ":"
               + second.ToString("D2");
    }
}
```

Class Factories in C#

Having created a class such as `TimeOfDayI`, we have an implementation and we can instantiate the `TimeOfDayI` class, but we cannot receive it as the return value or as an out-parameter from an operation invocation. To see why, consider the following simple interface:

Slice

```
interface Time {
    TimeOfDay get();
};
```

When a client invokes the `get` operation, the Ice run time must instantiate and return an instance of the `TimeOfDay` class. However, `TimeOfDay` is an abstract class that cannot be instantiated. Unless we tell it, the Ice run time cannot magically know that we have created a `TimeOfDayI` class that implements the abstract `format` operation of the `TimeOfDay` abstract class. In other words, we must provide the Ice run time with a factory that knows that the `TimeOfDay` abstract class has a `TimeOfDayI` concrete implementation. The `Ice::Communicator` interface provides us with the necessary operations:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local interface ObjectFactory {
        Object create(string type);
        void destroy();
    };

    local interface Communicator {
        void addObjectFactory(ObjectFactory factory, string id);
        ObjectFactory findObjectFactory(string id);
        // ...
    };
};
```

To supply the Ice run time with a factory for our `TimeOfDayI` class, we must implement the `ObjectFactory` interface:

C#

```
class ObjectFactory : Ice.ObjectFactory
{
    public Ice.Object create(string type)
    {
        if (type.Equals(M.TimeOfDay.ice_staticId()))
            return new TimeOfDayI();
        System.Diagnostics.Debug.Assert(false);
        return null;
    }

    public void destroy()
    {
        // Nothing to do
    }
}
```

The object factory's `create` method is called by the Ice run time when it needs to instantiate a `TimeOfDay` class. The factory's `destroy` m ethod is called by the Ice run time when its communicator is destroyed.

The `create` method is passed the type ID of the class to instantiate. For our `TimeOfDay` class, the type ID is `"::M::TimeOfDay"`. Our implementation of `create` checks the type ID: if it matches, the method instantiates and returns a `TimeOfDayI` object. For other type IDs, the method asserts because it does not know how to instantiate other types of objects.

Note that we used the `ice_staticId` method to obtain the type ID rather than embedding a literal string. Using a literal type ID string in your code is discouraged because it can lead to errors that are only detected at run time. For example, if a Slice class or one of its enclosing modules is renamed and the literal string is not changed accordingly, a receiver will fail to unmarshal the object and the Ice run time will raise `NoObjectFactoryException`. By using `ice_staticId` instead, we avoid any risk of a misspelled or obsolete type ID, and we can discover at compile time if a Slice class or module has been renamed.

Given a factory implementation, such as our `ObjectFactory`, we must inform the Ice run time of the existence of the factory:

Copyright 2016, ZeroC, Inc.

C#

```
Ice.Communicator ic = ...;
ic.addObjectFactory(new ObjectFactory(), M.TimeOfDay.ice_staticId());
```

Now, whenever the Ice run time needs to instantiate a class with the type ID `"::M::TimeOfDay"`, it calls the `create` method of the registered `ObjectFactory` instance, which returns a `TimeOfDayI` instance to the Ice run time.

The `destroy` operation of the object factory is invoked by the Ice run time when the communicator is destroyed. This gives you a chance to clean up any resources that may be used by your factory. Do not call `destroy` on the factory while it is registered with the communicator — if you do, the Ice run time has no idea that this has happened and, depending on what your `destroy` implementation is doing, may cause undefined behavior when the Ice run time tries to next use the factory.

The run time guarantees that `destroy` will be the last call made on the factory, that is, `create` will not be called concurrently with `destroy`, and `create` will not be called once `destroy` has been called. However, calls to `create` can be made concurrently.

Note that you cannot register a factory for the same type ID twice: if you call `addObjectFactory` with a type ID for which a factory is registered, the Ice run time throws an `AlreadyRegisteredException`.

Finally, keep in mind that if a class has only data members, but no operations, you need not create and register an object factory to transmit instances of such a class. Only if a class has operations do you have to define and register an object factory.

Class Constructors in C#

A class has a default constructor that initializes data members as follows:

Data Member Type	Default Value
`string`	Empty string
`enum`	Zero
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	Null
`dictionary`	Null
`class/interface`	Null

The constructor won't explicitly initialize a data member if the default C# behavior for that type produces the desired results.

If you wish to ensure that data members of primitive and enumerated types are initialized to specific values, you can declare default values in your Slice definition. The default constructor initializes each of these data members to its declared value instead.

A class also provides a constructor that accepts one argument for each member of the class. This allows you to create and initialize a class in a single statement, for example:

C#

```
TimeOfDayI tod = new TimeOfDayI(14, 45, 00); // 2:45pm
```

For a derived class, the constructor requires one argument for every member of the class, including inherited members. For example, consider the the definition from Class Inheritance once more:

Copyright 2016, ZeroC, Inc.

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
};

class DateTime extends TimeOfDay {
    short day;          // 1 - 31
    short month;        // 1 - 12
    short year;         // 1753 onwards
};
```

The Slice compiler generates the following constructors for these types:

Copyright 2016, ZeroC, Inc.

C#

```
public partial class TimeOfDay : Ice.ObjectImpl
{
    public TimeOfDay() {}

    public TimeOfDay(short hour, short minute, short second)
    {
        this.hour = hour;
        this.minute = minute;
        this.second = second;
    }

    // ...
}

public partial class DateTime : TimeOfDay
{
    public DateTime() : base() {}

    public DateTime(short hour,
                    short minute,
                    short second,
                    short day,
                    short month,
                    short year) : base(hour, minute, second)
    {
        this.day = day;
        this.month = month;
        this.year = year;
    }

    // ...
}
```

If you want to instantiate and initialize a `DateTime` instance, you must either use the default constructor or provide values for all of the data members of the instance, including data members of its base class.

For each optional data member of a class, the constructor accepts an `Ice.Optional` parameter of the appropriate type. You can pass the value `Ice.Util.None` as the value of any optional data member to initialize it to an unset condition.

See Also

- Classes
- Class Inheritance
- Type IDs
- C-Sharp Mapping for Structures
- C-Sharp Mapping for Optional Values
- The Current Object
- Dispatch Interceptors

Copyright 2016, ZeroC, Inc.

C-Sharp Mapping for Optional Values

On this page:

- The Ice.Optional Type
- The Ice.Util.None Value

The `Ice.Optional` Type

The C# mapping uses a generic structure to hold the values of optional data members and parameters:

C#

```
namespace Ice
{
    public struct NoneType {}

    [Serializable]
    public struct Optional<T> : System.ISerializable {
        public Optional(NoneType none);
        public Optional(T v);
        public Optional(Optional<T> v);

        public static explicit operator T(Optional<T> v);
        public static implicit operator Optional<T>(T v);
        public static implicit operator Optional<T>(NoneType v);

        public T Value { get; }
        public bool HasValue { get; }

        public override bool Equals(object other);
        public override int GetHashCode();

        ...
    }
}
```

The `Ice.Optional` type provides constructors and conversion operators that allow you to initialize an instance using the element type or an existing optional value. The default constructor initializes an instance to an unset condition. The `Value` property and conversion operator retrieve the current value held by the instance, or throw `System.InvalidOperationException` if no value is currently set. Use the `HasValue` property to test whether the instance has a value prior to accessing it.

The implicit conversion operators allow you to initialize an `Optional` instance without a cast. However, the conversion operator that extracts the current value is declared as `explicit` and therefore requires a cast:

C#

```
Ice.Optional<int> i = 5; // No cast necessary
int j = i;               // Error!
int k = (int)i;          // OK
```


Copyright 2016, ZeroC, Inc.

This operator raises an exception if no value is set, so the proper way to write this is shown below:

C#

```
Ice.Optional<int> i = 5;
int j;
if(i.HasValue)
    j = (int)i;
else
    ...
```

The `Ice.Util.None` Value

The `Ice.Optional` type provides a constructor and conversion operator that accept `NoneType`. Ice defines an instance of this type, `Ice.Util.None`, that you can use to initialize (or reset) an `Optional` instance to an unset condition:

C#

```
Ice.Optional<int> i = 5;
i = Ice.Util.None;
Debug.Assert(!i.HasValue); // true
```

You can pass `Ice.Util.None` anywhere an `Ice.Optional` value is expected.

See Also

- Optional Data Members

Copyright 2016, ZeroC, Inc.

Serializable Objects in C-Sharp

In addition to serializing Slice types, applications may also need to incorporate foreign types into their Slice definitions. Ice allows you to pass CLR serializable objects directly as operation parameters or as fields of another data type. For example:

Slice

```
["clr:serializable:SomeNamespace.CLRClass"]
sequence<byte> CLRObj;
struct MyStruct {
    int i;
    CLRObj o;
};

interface Example {
    void op(CLRObj o, MyStruct s);
};
```

The generated code for `MyStruct` contains member `i` of type `int` and a member `o` of type `SomeNamespace.CLRClass`:

C#

```
public partial class MyStruct : _System.ICloneable {
    public int i;
    public SomeNamespace.CLRClass o;

    // ...
}
```

Similarly, the signature for `op` has parameters of type `CLRClass` and `MyStruct`:

C#

```
void op(SomeNamespace.CLRClass o, MyStruct s);
```

Of course, your client and server code must have an implementation of `CLRClass` that sets the `Serializable` attribute:

Copyright 2016, ZeroC, Inc.

C#

```
namespace SomeNamespace {
    [Serializable]
    public class CLRClass {
        // ...
    }
}
```

You can implement this class in any way you see fit — the Ice run time does not place any other requirements on the implementation. However, note that the CLR requires the class to reside in the same assembly for client and server.

See Also

- Serializable Objects

Copyright 2016, ZeroC, Inc.

C-Sharp Specific Metadata Directives

The `slice2cs` compiler supports metadata directives that allow you to inject C# attribute specifications into the generated code. The metadata directive is `cs:attribute:`. For example:

Slice

```
["cs:attribute:System.Serializable"]
struct Point {
    double x;
    double y;
};
```

This results in the following code being generated for `S`:

C#

```
[System.Serializable]
public partial struct Point
{
    public double x;
    public double y;
    // ...
}
```

You can apply this metadata directive to any Slice construct, such as structure, operation, or parameter definitions.

You can use this directive also at global level. For example:

Slice

```
[["cs:attribute:assembly: AssemblyDescription(\"My assembly\")"]]
```

This results in the following code being inserted after any `using` directives and before any definitions:

C#

```
[assembly: AssemblyDescription("My assembly")]
```

See Also

- Metadata
- Slice Metadata Directives

Copyright 2016, ZeroC, Inc.

Asynchronous Method Invocation (AMI) in C-Sharp

Asynchronous Method Invocation (AMI) is the term used to describe the client-side support for the asynchronous programming model. AMI supports both oneway and twoway requests, but unlike their synchronous counterparts, AMI requests never block the calling thread. When a client issues an AMI request, the Ice run time hands the message off to the local transport buffer or, if the buffer is currently full, queues the request for later delivery. The application can then continue its activities and poll or wait for completion of the invocation, or receive a callback when the invocation completes.

AMI is transparent to the server: there is no way for the server to tell whether a client sent a request synchronously or asynchronously.

On this page:

- Basic Asynchronous API in C#
  - Asynchronous Proxy Methods in C#
  - Asynchronous Exception Semantics in C#
- AsyncResult Interface in C#
- Polling for Completion in C#
- Generic Completion Callbacks in C#
  - Using Cookies for Generic Completion Callbacks in C#
- Type-Safe Completion Callbacks in C#
  - Using Cookies for Type-Safe Completion Callbacks in C#
- Asynchronous Oneway Invocations in C#
- Flow Control in C#
- Asynchronous Batch Requests in C#
- Concurrency Semantics for AMI in C#

Basic Asynchronous API in C#

Consider the following simple Slice definition:

Slice

```
module Demo {
    interface Employees {
        string getName(int number);
    };
};
```

Asynchronous Proxy Methods in C#

Besides the synchronous proxy methods, `slice2cs` generates the following asynchronous proxy methods:

C#

```
public interface EmployeesPrx : Ice.ObjectPrx {
    Ice.AsyncResult<Demo.Callback_Employees_getName>
    begin_getName(int number);

    Ice.AsyncResult<Demo.Callback_Employees_getName>
    begin_getName(int number,
                  _System.Collections.Generic.Dictionary<string, string> ctx__);

    string end_getName(Ice.AsyncResult r__);
}
```


Copyright 2016, ZeroC, Inc.

Two additional overloads of `begin_getName` are generated for use with generic completion callbacks.

As you can see, the single `getName` operation results in `begin_getName` and `end_getName` methods. (The `begin_` method is overloaded so you can pass a per-invocation context.)

- The `begin_getName` method sends (or queues) an invocation of `getName`. This method does not block the calling thread.
- The `end_getName` method collects the result of the asynchronous invocation. If, at the time the calling thread calls `end_getName`, the result is not yet available, the calling thread blocks until the invocation completes. Otherwise, if the invocation completed some time before the call to `end_getName`, the method returns immediately with the result.

A client could call these methods as follows:

C#

```
EmployeesPrx e = ...;
Ice.AsyncResult r = e.begin_getName(99);

// Continue to do other things here...

string name = e.end_getName(r);
```

Because `begin_getName` does not block, the calling thread can do other things while the operation is in progress.

Note that `begin_getName` returns a value of type `Ice.AsyncResult`. (The class derives from `System.IAsyncResult`.) This value contains the state that the Ice run time requires to keep track of the asynchronous invocation. You must pass the `AsyncResult` that is returned by the `begin_` method to the corresponding `end_` method.

The `begin_` method has one parameter for each in-parameter of the corresponding Slice operation. Similarly, the `end_` method has one out-parameter for each out-parameter of the corresponding Slice operation (plus the `AsyncResult` parameter). For example, consider the following operation:

Slice

```
double op(int inp1, string inp2, out bool outp1, out long outp2);
```

The `begin_op` and `end_op` methods have the following signature:

C#

```
Ice.AsyncResult<Demo.Callback_Employees_op> begin_op(int inp1, string inp2);

double end_op(out bool outp1, out long outp2, Ice.AsyncResult r__);
```

Asynchronous Exception Semantics in C#

If an invocation raises an exception, the exception is thrown by the `end_` method, even if the actual error condition for the exception was encountered during the `begin_` method ("on the way out"). The advantage of this behavior is that all exception handling is located with the code that calls the `end_` method (instead of being present twice, once where the `begin_` method is called, and again where the `end_` meth od is called).

There is one exception to the above rule: if you destroy the communicator and then make an asynchronous invocation, the `begin_` method

Copyright 2016, ZeroC, Inc.

throws `CommunicatorDestroyedException`. This is necessary because, once the run time is finalized, it can no longer throw an exception from the `end_` method.

The only other exception that is thrown by the `begin_` and `end_` methods is `System.ArgumentException`. This exception indicates that you have used the API incorrectly. For example, the `begin_` method throws this exception if you call an operation that has a return value or out-parameters on a oneway proxy. Similarly, the `end_` method throws this exception if you use a different proxy to call the `end_` method than the proxy you used to call the `begin_` method, or if the `AsyncResult` you pass to the `end_` method was obtained by calling the `begin_` method for a different operation.

`AsyncResult` Interface in C#

The `AsyncResult` that is returned by the `begin_` method encapsulates the state of the asynchronous invocation:

C#

```
public interface AsyncResult : System.IAsyncResult
{
    void cancel();

    Ice.Communicator getCommunicator();
    Ice.Connection getConnection();
    ObjectPrx getProxy();
    string getOperation();
    object AsyncState { get; }

    bool IsCompleted { get; }
    void waitForCompleted();

    bool isSent();
    void waitForSent();

    void throwLocalException();

    bool sentSynchronously();

    AsyncResult whenSent(Ice.AsyncCallback cb);
    AsyncResult whenSent(Ice.SentCallback cb);
    AsyncResult whenCompleted(Ice.ExceptionCallback ex);
}

public interface AsyncResult<T> : AsyncResult
{
    AsyncResult<T> whenCompleted(T cb, Ice.ExceptionCallback excb);

    new AsyncResult<T> whenCompleted(Ice.ExceptionCallback excb);
    new AsyncResult<T> whenSent(Ice.SentCallback cb);
}
```

The methods and properties have the following semantics:

- `void cancel()`
  This method prevents a queued invocation from being sent or, if the invocation has already been sent, ignores a reply if the server sends one. `cancel` is a local operation and has no effect on the server. A canceled invocation is considered to be completed, meaning `IsCompleted` returns true, and the result of the invocation is an `Ice.InvocationCanceledException`.

Copyright 2016, ZeroC, Inc.

- `Communicator getCommunicator()`
  This method returns the communicator that sent the invocation.
- `Connection getConnection()`
  This method returns the connection that was used for the invocation. Note that, for typical asynchronous proxy invocations, this method returns a nil value because the possibility of automatic retries means the connection that is currently in use could change unexpectedly. The `getConnection` method only returns a non-nil value when the `AsyncResult` object is obtained by calling `begin_flushBatchRequests` on a `Connection` object.
- `ObjectPrx getProxy()`
  This method returns the proxy that was used to call the `begin_` method, or nil if the `AsyncResult` object was not obtained via an asynchronous proxy invocation.
- `string getOperation()`
  This method returns the name of the operation.
- `object AsyncState { get; }`
  This property stores an object that you can use to pass shared state from the `begin_` to the `end_` method.
- `bool IsCompleted { get; }`
  This property is true if, at the time it is called, the result of an invocation is available, indicating that a call to the `end_` method will not block the caller. Otherwise, if the result is not yet available, the method returns false.
- `void waitForCompleted()`
  This method blocks the caller until the result of an invocation becomes available.
- `bool isSent()`
  When you call the `begin_` method, the Ice run time attempts to write the corresponding request to the client-side transport. If the transport cannot accept the request, the Ice run time queues the request for later transmission. `isSent` returns true if, at the time it is called, the request has been written to the local transport (whether it was initially queued or not). Otherwise, if the request is still queued or an exception occurred before the request could be sent, `isSent` returns false.
- `void waitForSent()`
  This method blocks the calling thread until a request has been written to the client-side transport, or an exception occurs. After `waitForSent` returns, `isSent` returns true if the request was successfully written to the client-side transport, or false if an exception occurred. In the case of a failure, you can call the corresponding `end_` method or `throwLocalException` to obtain the exception.
- `void throwLocalException()`
  This method throws the local exception that caused the invocation to fail. If no exception has occurred yet, `throwLocalException` does nothing.
- `bool sentSynchronously()`
  This method returns true if a request was written to the client-side transport without first being queued. If the request was initially queued, `sentSynchronously` returns false (independent of whether the request is still in the queue or has since been written to the client-side transport).
- `AsyncResult whenSent(Ice.SentCallback cb)`
  `AsyncResult<T> whenSent(Ice.SentCallback cb)`
  `AsyncResult whenCompleted(Ice.ExceptionCallback ex)`
  `AsyncResult<T> whenCompleted(T cb,Ice.ExceptionCallback excb)`
  `AsyncResult<T> whenCompleted(Ice.ExceptionCallback excb)`
  These methods allow you to specify callback methods that are called by the Ice run time. The `whenSent` methods set a callback that triggers when an asynchronous invocation is written to the client-side transport. The `whenCompleted` methods set a callback that triggers when an asynchronous invocation completes (also see Generic Completion Callbacks in C#).

Polling for Completion in C#

The `AsyncResult` methods allow you to poll for call completion. Polling is useful in a variety of cases. As an example, consider the following simple interface to transfer files from client to server:

Copyright 2016, ZeroC, Inc.

Slice

```
interface FileTransfer
{
    void send(int offset, ByteSeq bytes);
};
```

The client repeatedly calls `send` to send a chunk of the file, indicating at which offset in the file the chunk belongs. A naïve way to transmit a file would be along the following lines:

C#

```
FileHandle file = open(...);
FileTransferPrx ft = ...;
const int chunkSize = ...;
int offset = 0;
while (!file.eof()) {
    byte[] bs;
    bs = file.read(chunkSize); // Read a chunk
    ft.send(offset, bs);       // Send the chunk
    offset += bs.Length;
}
```

This works, but not very well: because the client makes synchronous calls, it writes each chunk on the wire and then waits for the server to receive the data, process it, and return a reply before writing the next chunk. This means that both client and server spend much of their time doing nothing — the client does nothing while the server processes the data, and the server does nothing while it waits for the client to send the next chunk.

Using asynchronous calls, we can improve on this considerably:

Copyright 2016, ZeroC, Inc.

C#

```
FileHandle file = open(...);
FileTransferPrx ft = ...;
const int chunkSize = ...;
int offset = 0;

LinkedList<Ice.AsyncResult> results = new LinkedList<Ice.AsyncResult>();
const int numRequests = 5;

while (!file.eof()) {
    byte[] bs;
    bs = file.read(chunkSize);

    // Send up to numRequests + 1 chunks asynchronously.
    Ice.AsyncResult r = ft.begin_send(offset, bs);
    offset += bs.Length;

    // Wait until this request has been passed to the transport.
    r.waitForSent();
    results.AddLast(r);

    // Once there are more than numRequests, wait for the least
    // recent one to complete.
    while (results.Count > numRequests) {
        Ice.AsyncResult r = results.First;
        results.RemoveFirst();
        r.waitForCompleted();
    }
}

// Wait for any remaining requests to complete.
while (results.Count > 0) {
    Ice.AsyncResult r = results.First;
    results.RemoveFirst();
    r.waitForCompleted();
}
```

With this code, the client sends up to `numRequests + 1` chunks before it waits for the least recent one of these requests to complete. In other words, the client sends the next request without waiting for the preceding request to complete, up to the limit set by `numRequests`. In effect, this allows the client to "keep the pipe to the server full of data": the client keeps sending data, so both client and server continuously do work.

Obviously, the correct chunk size and value of `numRequests` depend on the bandwidth of the network as well as the amount of time taken by the server to process each request. However, with a little testing, you can quickly zoom in on the point where making the requests larger or queuing more requests no longer improves performance. With this technique, you can realize the full bandwidth of the link to within a percent or two of the theoretical bandwidth limit of a native socket connection.

Generic Completion Callbacks in C#

The `begin_` method is overloaded to allow you to provide completion callbacks. Here are the corresponding methods for the `getName` operation:

Copyright 2016, ZeroC, Inc.

C#

```
Ice.AsyncResult begin_getName(
        int number,
        Ice.AsyncCallback cb__,
        object cookie__);

Ice.AsyncResult begin_getName(
        int number,
        _System.Collections.Generic.Dictionary<string, string> ctx__,
        Ice.AsyncCallback cb__,
        object cookie__);
```

The second version of `begin_getName` lets you override the default context. (We discuss the purpose of the `cookie` parameter in the next section.) Following the in-parameters, the `begin_` method accepts a parameter of type `Ice.AsyncCallback`, which is a delegate for a callback method. The Ice run time invokes the callback method when an asynchronous operation completes. Your callback method must have `void` return type and accept a single parameter of type `AsyncResult`, for example:

C#

```
private class MyCallback
{
    public void finished(Ice.AsyncResult r)
    {
        EmployeesPrx e = (EmployeesPrx)r.getProxy();
        try {
            string name = e.end_getName(r);
            System.Console.WriteLine("Name is: " + name);
        } catch (Ice.Exception ex) {
            System.Console.Err.WriteLine("Exception is: " + ex);
        }
    }
}
```

The implementation of your callback method must call the `end_` method. The proxy for the call is available via the `getProxy` method on the `AsyncResult` that is passed by the Ice run time. The return type of `getProxy` is `Ice.ObjectPrx`, so you must down-cast the proxy to its correct type.

Your callback method should catch and handle any exceptions that may be thrown by the `end_` method. If you allow an exception to escape from the callback method, the Ice run time produces a log entry by default and ignores the exception. (You can disable the log message by setting the property `Ice.Warn.AMICallback` to zero.)

To inform the Ice run time that you want to receive a callback for the completion of the asynchronous call, you pass a delegate for your callback method to the `begin_` method:

Copyright 2016, ZeroC, Inc.

C#

```
EmployeesPrx e = ...;

MyCallback cb = new MyCallback();
Ice.AsyncCallback del = new Ice.AsyncCallback(cb.finished);

e.begin_getName(99, del, null);
```

The trailing `null` argument specifies a cookie, which we will discuss shortly.

You can avoid explicit instantiation of the delegate and, more tersely, write:

C#

```
EmployeesPrx e = ...;

MyCallback cb = new MyCallback();
e.begin_getName(99, cb.finished, null);
```

In fact, there's often no need to define a callback class at all. An anonymous method is one alternative:

C#

```
EmployeesPrx e = ...;

e.begin_getName(99,
    delegate(Ice.AsyncResult r)
    {
        EmployeesPrx e = (EmployeesPrx)r.getProxy();
        try {
            string name = e.end_getName(r);
            System.Console.WriteLine("Name is: " + name);
        } catch (Ice.Exception ex) {
            System.Console.Err.WriteLine("Exception is: " + ex);
        }
    }, null);
```

A lambda is another alternative:

Copyright 2016, ZeroC, Inc.

C#

```
EmployeesPrx e = ...;

e.begin_getName(99,
    r => {
        EmployeesPrx e = (EmployeesPrx)r.getProxy();
        try {
            string name = e.end_getName(r);
            System.Console.WriteLine("Name is: " + name);
        } catch (Ice.Exception ex) {
            System.Console.Err.WriteLine("Exception is: " + ex);
        }
    }, null);
```

Using Cookies for Generic Completion Callbacks in C#

It is common for the `end_` method to require access to some state that is established by the code that calls the `begin_` method. As an example, consider an application that asynchronously starts a number of operations and, as each operation completes, needs to update different user interface elements with the results. In this case, the `begin_` method knows which user interface element should receive the update, and the `end_` method needs access to that element.

The API allows you to pass such state by providing a cookie. A cookie is any class instance; the class can contain whatever data you want to pass, as well as any methods you may want to add to manipulate that data.

Here is an example implementation that stores a `Widget`. (We assume that this class provides whatever methods are needed by the `end_` method to update the display.) When you call the `begin_` method, you pass the appropriate cookie instance to inform the `end_` method how to update the display:

C#

```
// Invoke the getName operation with different widget cookies.
MyCallback cb = ...;
e.begin_getName(99, cb.finished, widget1);
e.begin_getName(24, cb.finished, widget2);
```

The `end_` method can retrieve the cookie from the `AsyncResult` by reading the `AsyncState` property. For this example, we assume that widgets have a `writeString` method that updates the relevant UI element:

Copyright 2016, ZeroC, Inc.

C#

```
public void finished(Ice.AsyncResult r)
{
    EmployeesPrx e = (EmployeesPrx)r.getProxy();
    Widget widget = (Widget)r.AsyncState;
    try {
        string name = e.end_getName(r);
        widget.writeString(name);
    } catch (Ice.Exception ex) {
        handleException(ex);
    }
}
```

The cookie provides a simple and effective way for you to pass state between the point where an operation is invoked and the point where its results are processed. Moreover, if you have a number of operations that share common state, you can pass the same cookie instance to multiple invocations.

Type-Safe Completion Callbacks in C#

The generic callback API is not entirely type-safe:

- You must down-cast the return value of `getProxy` to the correct proxy type before you can call the `end_` method.
- You must call the correct `end_` method to match the operation called by the `begin_` method.
- You must remember to catch exceptions when you call the `end_` method; if you forget to do this, you will not know that the operation failed.

`slice2cs` generates an additional type-safe API that takes care of these chores for you. To use type-safe callbacks, you supply delegates for two callback methods:

- a success callback that is called if the operation succeeds
- a failure callback that is called if the operation raises an exception

Here is a callback class for an invocation of the `getName` operation:

C#

```
public class MyCallback
{
    public void getNameCB(string name)
    {
        System.Console.WriteLine("Name is: " + name);
    }

    public void failureCB(Ice.Exception ex)
    {
        System.Console.Err.WriteLine("Exception is: " + ex);
    }
}
```

The callback methods can have any name you prefer and must have `void` return type. The failure callback always has a single parameter of

Copyright 2016, ZeroC, Inc.

type `Ice.Exception`. The success callback parameters depend on the operation signature. If the operation has non-`void` return type, the first parameter of the success callback is the return value. The return value (if any) is followed by a parameter for each out-parameter of the corresponding Slice operation, in the order of declaration.

At the calling end, you call the `begin_` method as follows:

C#

```
MyCallback cb = new MyCallback();

e.begin_getName(99).whenCompleted(cb.getNameCB, cb.failureCB);
```

Note the `whenCompleted` method on the `AsyncResult` that is returned by the `begin_` method. This method establishes the link between the `begin_` method and the callbacks that are called by the Ice run time by setting the delegates for the success and failure methods.

It is legal to pass a null delegate for the success or failure methods. For the success callback, this is legal only for operations that have `void` return type and no out-parameters. This is useful if you do not care when the operation completes but want to know if the call failed. If you pass a null exception delegate, the Ice run time will ignore any exception that is raised by the invocation.

Defining a callback class as we've shown above is only necessary in practice when the callback has additional state to manage. In many cases, there's no need for a callback class and the delegates can be defined inline:

C#

```
e.begin_getName(99).whenCompleted(
    delegate(string name)
    {
        System.Console.WriteLine("Name is: " + name);
    },
    delegate(Ice.Exception ex)
    {
        System.Console.Err.WriteLine("Exception is: " + ex);
    });
```

Using lambda functions makes the callbacks even more compact:

C#

```
e.begin_getName(99).whenCompleted(
    name => {
        System.Console.WriteLine("Name is: " + name);
    },
    ex => {
        System.Console.Err.WriteLine("Exception is: " + ex);
    });
```

Using Cookies for Type-Safe Completion Callbacks in C#

Copyright 2016, ZeroC, Inc.

The type-safe API does not support cookies. If you want to pass state from the `begin_` method to the `end_` method, you must use the generic API or, alternatively, place the state into the callback class containing the callback methods. Here is a simple implementation of a callback class that stores a widget that can be retrieved by the `end_` method:

C#

```
public class MyCallback
{
    public MyCallback(Widget w)
    {
        _w = w;
    }

    private Widget _w;

    public void getNameCB(string name)
    {
        _w.writeString(name);
    }

    public void failureCB(Ice.Exception ex)
    {
        _w.writeError(ex);
    }
}
```

When you call the `begin_` method, you pass the appropriate callback instance to inform the `end_` method how to update the display:

C#

```
EmployeesPrx e = ...;
Widget widget1 = ...;
Widget widget2 = ...;

MyCallback cb1 = new MyCallback(widget1);
MyCallback cb2 = new MyCallback(widget2);

// Invoke the getName operation with different widget callbacks.

e.begin_getName(99).whenCompleted(cb1.getNameCB, cb1.failureCB);
e.begin_getName(24).whenCompleted(cb2.getNameCB, cb2.failureCB);
```

Asynchronous Oneway Invocations in C#

You can invoke operations via oneway proxies asynchronously, provided the operation has `void` return type, does not have any out-parameters, and does not raise user exceptions. If you call the `begin_` method on a oneway proxy for an operation that returns values or raises a user exception, the `begin_` method throws a `System.ArgumentException`.

For the generic API, the callback method looks exactly as for a twoway invocation. However, for oneway invocations, the Ice run time does not call the callback method unless the invocation raised an exception during the `begin_` method ("on the way out").

Copyright 2016, ZeroC, Inc.

For the type-safe API, you only specify a delegate for the failure method. For example, here is how you could call `ice_ping` asynchronously :

C#

```
ObjectPrx p = ...;
MyCallback cb = new MyCallback();
p.begin_ice_ping().whenCompleted(cb.failureCB);
```

Flow Control in C#

Asynchronous method invocations never block the thread that calls the `begin_` method: the Ice run time checks to see whether it can write the request to the local transport. If it can, it does so immediately in the caller's thread. (In that case, `AsyncResult.sentSynchronously` returns true.) Alternatively, if the local transport does not have sufficient buffer space to accept the request, the Ice run time queues the request internally for later transmission in the background. (In that case, `AsyncResult.sentSynchronously` returns false.)

This creates a potential problem: if a client sends many asynchronous requests at the time the server is too busy to keep up with them, the requests pile up in the client-side run time until, eventually, the client runs out of memory.

The API provides a way for you to implement flow control by counting the number of requests that are queued so, if that number exceeds some threshold, the client stops invoking more operations until some of the queued operations have drained out of the local transport.

For the generic API, you can create an additional callback method:

C#

```
public class MyCallback
{
    public void finished(Ice.AsyncResult r)
    {
        // ...
    }

    public void sent(Ice.AsyncResult r)
    {
        // ...
    }
}
```

As with any other callback method, you are free to choose any name you like. For this example, the name of the callback method is `sent`. You inform the Ice run time that you want to be informed when a call has been passed to the local transport by calling `whenSent`:

C#

```
MyCallback cb = new MyCallback();

e.begin_getName(99).whenCompleted(cb.getNameCB, cb.failureCB).whenSent(cb.sent);
```


Copyright 2016, ZeroC, Inc.

If the Ice run time can immediately pass the request to the local transport, it does so and invokes the `sent` method from the thread that calls the `begin_` method. On the other hand, if the run time has to queue the request, it calls the `sent` method from a different thread once it has written the request to the local transport. In addition, you can find out from the `AsyncResult` that is returned by the `begin_` method whether the request was sent synchronously or was queued, by calling `sentSynchronously`.

For the generic API, the `sent` method has the following signature:

C#

```
void sent(Ice.AsyncResult r);
```

For the type-safe API, the signature is:

C#

```
void sent(bool sentSynchronously);
```

For the generic API, you can find out whether the request was sent synchronously by calling `sentSynchronously` on the `AsyncResult`. For the type-safe API, the boolean `sentSynchronously` parameter provides the same information.

The `sent` methods allow you to limit the number of queued requests by counting the number of requests that are queued and decrementing the count when the Ice run time passes a request to the local transport.

Asynchronous Batch Requests in C#

Applications that send batched requests can either flush a batch explicitly or allow the Ice run time to flush automatically. The proxy method `ice_flushBatchRequests` performs an immediate flush using the synchronous invocation model and may block the calling thread until the entire message can be sent. Ice also provides asynchronous versions of this method so you can flush batch requests asynchronously.

`begin_ice_flushBatchRequests` and `end_ice_flushBatchRequests` are proxy methods that flush any batch requests queued by that proxy.

In addition, similar methods are available on the communicator and the `Connection` object that is returned by `AsyncResult.getConnection`. These methods flush batch requests sent via the same communicator and via the same connection, respectively.

Concurrency Semantics for AMI in C#

The Ice run time always invokes your callback methods from a separate thread, with one exception: it calls the `sent` callback from the thread calling the `begin_` method if the request could be sent synchronously. In the `sent` callback, you know which thread is calling the callback by looking at the `sentSynchronously` member or parameter.

See Also

- Request Contexts
- Batched Invocations
- Collocated Invocation and Dispatch

Copyright 2016, ZeroC, Inc.

slice2cs Command-Line Options

The Slice-to-C# compiler, `slice2cs`, offers the following command-line options in addition to the standard options described in Using the Slice Compilers:

- `--tie`
  Generate tie classes.
- `--impl`
  Generate sample implementation files. This option will not overwrite an existing file.
- `--impl-tie`
  Generate sample implementation files using tie classes. This option will not overwrite an existing file.
- `--checksum`
  Generate checksums for Slice definitions.
- `--stream`
  Generate streaming helper methods for Slice types.

See Also

- Using the Slice Compilers
- Tie Classes in C-Sharp
- Streaming Interfaces

Copyright 2016, ZeroC, Inc.

Using Slice Checksums in C-Sharp

The Slice compilers can optionally generate checksums of Slice definitions. For `slice2cs`, the `--checksum` option causes the compiler to generate checksums in each C# source file that are added to a member of the `Ice.SliceChecksums` class:

C#

```
namespace Ice {
    public sealed class SliceChecksums {
        public readonly static System.Collections.Generic.Dictionary<string,
string> checksums;
    }
}
```

The `checksums` map is initialized automatically prior to first use; no action is required by the application.

In order to verify a server's checksums, a client could simply compare the dictionaries using the `Equals` function. However, this is not feasible if it is possible that the server might be linked with more Slice definitions than the client. A more general solution is to iterate over the local checksums as demonstrated below:

C#

```
System.Collections.Generic.Dictionary<string,string> serverChecksums = ...
foreach(System.Collections.Generic.KeyValuePair<string, string> e
in Ice.SliceChecksums.checksums) {
    string checksum;
    if (!serverChecksums.TryGetValue(e.Key, out checksum)) {
        // No match found for type id!
    } else if (!checksum.Equals(e.Value)) {
        // Checksum mismatch!
    }
}
```

In this example, the client first verifies that the server's dictionary contains an entry for each Slice type ID, and then it proceeds to compare the checksums.

See Also

- Slice Checksums

Copyright 2016, ZeroC, Inc.

Example of a File System Client in C-Sharp

This page presents a very simple client to access a server that implements the file system we developed in Slice for a Simple File System. The C# code hardly differs from the code you would write for an ordinary C# program. This is one of the biggest advantages of using Ice: accessing a remote object is as easy as accessing an ordinary, local C# object. This allows you to put your effort where you should, namely, into developing your application logic instead of having to struggle with arcane networking APIs. This is true for the server side as well, meaning that you can develop distributed applications easily and efficiently.

We now have seen enough of the client-side C# mapping to develop a complete client to access our remote file system. For reference, here is the Slice definition once more:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };

    exception GenericError {
        string reason;
    };

    sequence<string> Lines;

    interface File extends Node {
        idempotent Lines read();
        idempotent void write(Lines text) throws GenericError;
    };

    sequence<Node*> NodeSeq;

    interface Directory extends Node {
        idempotent NodeSeq list();
    };
};
```

To exercise the file system, the client does a recursive listing of the file system, starting at the root directory. For each node in the file system, the client shows the name of the node and whether that node is a file or directory. If the node is a file, the client retrieves the contents of the file and prints them.

The body of the client code looks as follows:

C#

```
using System;
using Filesystem;

public class Client
{
    // Recursively print the contents of directory "dir"
    // in tree fashion. For files, show the contents of
    // each file. The "depth" parameter is the current
    // nesting level (for indentation).
```


Copyright 2016, ZeroC, Inc.

```
    static void listRecursive(DirectoryPrx dir, int depth)
    {
        string indent = new string('\t', ++depth);

        NodePrx[] contents = dir.list();

        foreach (NodePrx node in contents)
            DirectoryPrx subdir = DirectoryPrxHelper.checkedCast(node);
            FilePrx file = FilePrxHelper.uncheckedCast(node);
            Console.WriteLine(
                indent + node.name() +
(subdir != null ? " (directory):" : " (file):"));
            if (subdir != null) {
                listRecursive(subdir, depth);
            } else {
                string[] text = file.read();
                for (int j = 0; j < text.Length; ++j)
                    Console.WriteLine(indent + "\t" + text[j]);
            }
        }
    }

    public static void Main(string[] args)
    {
        int status = 0;
        try {
            // Create a communicator
            //
            using(Ice.Communicator ic = Ice.Util.initialize(ref args)) {
                // Create a proxy for the root directory
                //
                Ice.ObjectPrx obj = ic.stringToProxy("RootDir:default -p 10000");

                // Down-cast the proxy to a Directory proxy
                //
                DirectoryPrx rootDir = DirectoryPrxHelper.checkedCast(obj);

                // Recursively list the contents of the root directory
                //
                Console.WriteLine("Contents of root directory:");
                listRecursive(rootDir, 0);
            }
        } catch (Exception e) {
            Console.Error.WriteLine(e);
            status = 1;
        }
        Environment.Exit(status);
```


Copyright 2016, ZeroC, Inc.

```
    }
}
```

The `Client` class defines two methods: `listRecursive`, which is a helper function to print the contents of the file system, and `Main`, which is the main program. Let us look at `Main` first:

1. The structure of the code in `Main` follows what we saw in Hello World Application. After initializing the run time, the client creates a proxy to the root directory of the file system. For this example, we assume that the server runs on the local host and listens using the default protocol (TCP/IP) at port 10000. The object identity of the root directory is known to be `RootDir`.
2. The client down-casts the proxy to `DirectoryPrx` and passes that proxy to `listRecursive`, which prints the contents of the file system.

Most of the work happens in `listRecursive`. The function is passed a proxy to a directory to list, and an indent level. (The indent level increments with each recursive call and allows the code to print the name of each node at an indent level that corresponds to the depth of the tree at that node.) `listRecursive` calls the `list` operation on the directory and iterates over the returned sequence of nodes:

1. The code does a `checkedCast` to narrow the `Node` proxy to a `Directory` proxy, as well as an `uncheckedCast` to narrow the `Node` proxy to a `File` proxy. Exactly one of those casts will succeed, so there is no need to call `checkedCast` twice: if the `Node` *is-a* `Directory`, the code uses the `DirectoryPrx` returned by the `checkedCast`; if the `checkedCast` fails, we *know* that the Node *is-a* File and, therefore, an `uncheckedCast` is sufficient to get a `FilePrx`.
   In general, if you know that a down-cast to a specific type will succeed, it is preferable to use an `uncheckedCast` instead of a `checkedCast` because an `uncheckedCast` does not incur any network traffic.
2. The code prints the name of the file or directory and then, depending on which cast succeeded, prints `"(directory)"` or `"(file)"` following the name.
3. The code checks the type of the node:
   - If it is a directory, the code recurses, incrementing the indent level.
   - If it is a file, the code calls the `read` operation on the file to retrieve the file contents and then iterates over the returned sequence of lines, printing each line.

Assume that we have a small file system consisting of two files and a directory as follows:

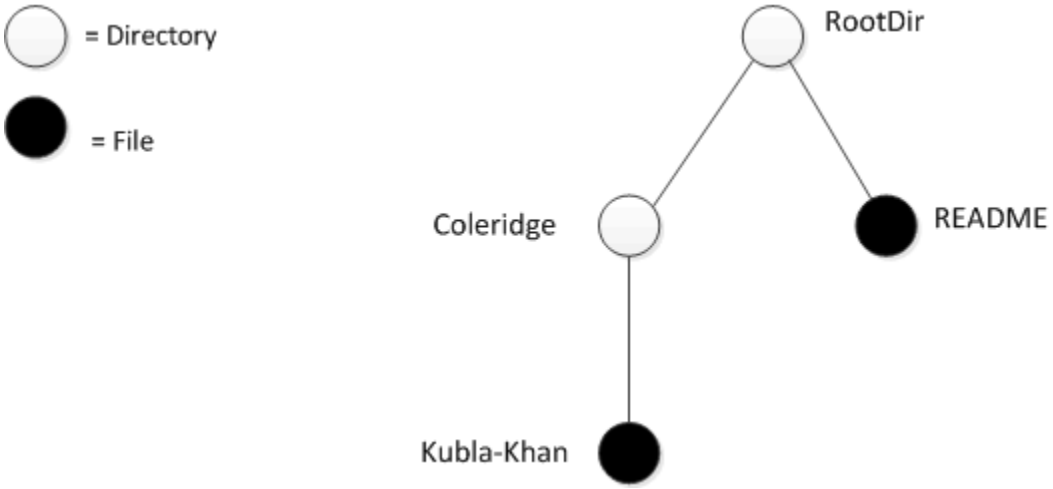

A small file system.

The output produced by the client for this file system is:

Copyright 2016, ZeroC, Inc.

```
Contents of root directory:
        README (file):
                This file system contains a collection of poetry.
        Coleridge (directory):
                Kubla_Khan (file):
                        In Xanadu did Kubla Khan
                        A stately pleasure-dome decree:
                        Where Alph, the sacred river, ran
                        Through caverns measureless to man
                        Down to a sunless sea.
```

Note that, so far, our client (and server) are not very sophisticated:

- The protocol and address information are hard-wired into the code.
- The client makes more remote procedure calls than strictly necessary; with minor redesign of the Slice definitions, many of these calls can be avoided.

We will see how to address these shortcomings in our discussions of IceGrid and object life cycle.

See Also

- Hello World Application
- Slice for a Simple File System
- Example of a File System Server in C-Sharp
- Object Life Cycle
- IceGrid

Copyright 2016, ZeroC, Inc.

Server-Side Slice-to-C-Sharp Mapping

The mapping for Slice data types to C# is identical on the client side and server side. This means that everything in the Client-Side Slice-to-C-Sharp Mapping also applies to the server side. However, for the server side, there are a few additional things you need to know — specifically, how to:

- Initialize and finalize the server-side run time
- Implement servants
- Pass parameters and throw exceptions
- Create servants and register them with the Ice run time.

Because the mapping for Slice data types is identical for clients and servers, the server-side mapping only adds a few additional mechanisms to the client side: a small API to initialize and finalize the run time, plus a few rules for how to derive servant classes from skeletons and how to register servants with the server-side run time.

Although the examples we present in this chapter are very simple, they accurately reflect the basics of writing an Ice server. Of course, for more sophisticated servers, you will be using additional APIs, for example, to improve performance or scalability. However, these APIs are all described in Slice, so, to use these APIs, you need not learn any C# mapping rules beyond those described here.

Topics

- The Server-Side main Method in C-Sharp
- Server-Side C-Sharp Mapping for Interfaces
- Parameter Passing in C-Sharp
- Raising Exceptions in C-Sharp
- Tie Classes in C-Sharp
- Object Incarnation in C-Sharp
- Asynchronous Method Dispatch (AMD) in C-Sharp
- Example of a File System Server in C-Sharp

Copyright 2016, ZeroC, Inc.

The Server-Side main Method in C-Sharp

On this page:

- A Basic Main Method in C#
- The Ice.Application Class in C#
  - Using Ice.Application on the Client Side in C#
  - Catching Signals in C#
  - Ice.Application and Properties in C#
  - Limitations of Ice.Application in C#

A Basic `Main` Method in C#

The main entry point to the Ice run time is represented by the local interface `Ice::Communicator`. As for the client side, you must initialize the Ice run time by calling `Ice.Util.initialize` before you can do anything else in your server. `Ice.Util.initialize` returns a reference to an instance of an `Ice.Communicator`:

C#

```
using System;

public class Server
{
    public static void Main(string[] args)
    {
        int status = 0;
        Ice.Communicator communicator = null;

        try {
            communicator = Ice.Util.initialize(ref args);
            // ...
        } catch (Exception ex) {
            Console.Error.WriteLine(ex);
            status = 1;
        }
        // ...
    }
}
```

`Ice.Util.initialize` accepts the argument vector that is passed to `Main` by the operating system. The method scans the argument vector for any command-line options that are relevant to the Ice run time; any such options are removed from the argument vector so, when `Ice.Util.initialize` returns, the only options and arguments remaining are those that concern your application. If anything goes wrong during initialization, `initialize` throws an exception.

Before leaving your `Main` method, you *must* call `Communicator.destroy`. The `destroy` operation is responsible for finalizing the Ice run time. In particular, `destroy` waits for any operation implementations that are still executing in the server to complete. In addition, `destroy` e nsures that any outstanding threads are joined with and reclaims a number of operating system resources, such as file descriptors and memory. Never allow your `Main` method to terminate without calling `destroy` first; doing so has undefined behavior.

The general shape of our server-side `Main` method is therefore as follows:

Copyright 2016, ZeroC, Inc.

C#

```
using System;

public class Server
{
    public static void Main(string[] args)
    {
        int status = 0;
        Ice.Communicator communicator = null;

        try {
            communicator = Ice.Util.initialize(ref args);
            // ...
        } catch (Exception ex) {
            Console.Error.WriteLine(ex);
            status = 1;
        }
        if (communicator != null) {
            try {
                communicator.destroy();
            } catch (Exception ex) {
                Console.Error.WriteLine(ex);
                status = 1;
            }
        }
        Environment.Exit(status);
    }
}
```

Note that the code places the call to `Ice.Util.initialize` into a `try` block and takes care to return the correct exit status to the operating system. Also note that an attempt to destroy the communicator is made only if the initialization succeeded.

A simpler version of the `Main` method employs a `using` statement to automatically dispose of the communicator:

Copyright 2016, ZeroC, Inc.

C#

```
using System;

public class Server
{
    public static void Main(string[] args)
    {
        int status = 0;

        try {
            using(Ice.Communicator communicator = Ice.Util.initialize(ref args)) {
                // ...
            }
        } catch (Exception ex) {
            Console.Error.WriteLine(ex);
            status = 1;
        }
        Environment.Exit(status);
    }
}
```

Communicator objects implement `IDisposable`, and calling `Dispose` on a communicator is equivalent to invoking `destroy`.

The `Ice.Application` Class in C#

The preceding structure for the `Main` method is so common that Ice offers a class, `Ice.Application`, that encapsulates all the correct initialization and finalization activities. The synopsis of the class is as follows (with some detail omitted for now):

Copyright 2016, ZeroC, Inc.

C#

```
namespace Ice
{
    public abstract class Application
    {
        public abstract int run(string[] args);

        public Application();

        public Application(SignalPolicy signalPolicy);

        public int main(string[] args);
        public int main(string[] args, string configFile);
        public int main(string[] args, InitializationData init);

        public static string appName();

        public static Communicator communicator();
    }
}
```

The intent of this class is that you specialize `Ice.Application` and implement the abstract `run` method in your derived class. Whatever code you would normally place in `Main` goes into the `run` method instead. Using `Ice.Application`, our program looks as follows:

C#

```
using System;

public class Server
{
    class App : Ice.Application
    {
        public override int run(string[] args)
        {
            // Server code here...

            return 0;
        }
    }

    public static void Main(string[] args)
    {
        App app = new App();
        Environment.Exit(app.main(args));
    }
}
```

Note that `Application.main` is overloaded: you can pass an optional file name or an `InitializationData` structure.

Copyright 2016, ZeroC, Inc.

If you pass a configuration file name to `main`, the property settings in this file are overridden by settings in a file identified by the `ICE_CONFIG` environment variable (if defined). Property settings supplied on the command line take precedence over all other settings.

The `Application.main` method does the following:

1. It installs an exception handler for `System.Exception`. If your code fails to handle an exception, `Application.main` prints the name of the exception and a stack trace on `Console.Error` before returning with a non-zero return value.
2. It initializes (by calling `Ice.Util.initialize`) and finalizes (by calling `Communicator.destroy`) a communicator. You can get access to the communicator for your server by calling the static `communicator` accessor.
3. It scans the argument vector for options that are relevant to the Ice run time and removes any such options. The argument vector that is passed to your `run` method therefore is free of Ice-related options and only contains options and arguments that are specific to your application.
4. It provides the name of your application via the static `appName` method. You can get at the application name from anywhere in your code by calling `Ice.Application.appName` (which is usually required for error messages).
5. It installs a signal handler that properly destroys the communicator.
6. It installs a per-process logger if the application has not already configured one. The per-process logger uses the value of the `Ice.ProgramName` property as a prefix for its messages and sends its output to the standard error channel. An application can also specify an alternate logger.

Using `Ice.Application` ensures that your program properly finalizes the Ice run time, whether your server terminates normally or in response to an exception. We recommend that all your programs use this class; doing so makes your life easier. In addition, `Ice.Application` also provides features for signal handling and configuration that you do not have to implement yourself when you use this class.

Using `Ice.Application` on the Client Side in C#

You can use `Ice.Application` for your clients as well: simply implement a class that derives from `Ice.Application` and place the client code into its `run` method. The advantage of this approach is the same as for the server side: `Ice.Application` ensures that the communicator is destroyed correctly even in the presence of exceptions.

Catching Signals in C#

The simple server we developed in Hello World Application had no way to shut down cleanly: we simply interrupted the server from the command line to force it to exit. Terminating a server in this fashion is unacceptable for many real-life server applications: typically, the server has to perform some cleanup work before terminating, such as flushing database buffers or closing network connections. This is particularly important on receipt of a signal or keyboard interrupt to prevent possible corruption of database files or other persistent data.

To make it easier to deal with signals, `Ice.Application` encapsulates the low-level signal handling tasks, allowing you to cleanly shut down on receipt of a signal.

C#

```
namespace Ice
{
    public abstract class Application
    {
        // ...

        public static void destroyOnInterrupt();
        public static void shutdownOnInterrupt();
        public static void ignoreInterrupt();
        public static void callbackOnInterrupt();
        public static void holdInterrupt();
        public static void releaseInterrupt();

        public static bool interrupted();

        public virtual void interruptCallback(int sig);
    }
}
```


Copyright 2016, ZeroC, Inc.

The methods behave as follows:

- `destroyOnInterrupt`
  This method installs a handler that destroys the communicator if it is interrupted. This is the default behavior.
- `shutdownOnInterrupt`
  This method installs a handler that shuts down the communicator if it is interrupted.
- `ignoreInterrupt`
  This method causes signals to be ignored.
- `callbackOnInterrupt`
  This method configures `Ice.Application` to invoke `interruptCallback` when a signal occurs, thereby giving the subclass responsibility for handling the signal.
- `holdInterrupt`
  This method temporarily blocks signal delivery.
- `releaseInterrupt`
  This method restores signal delivery to the previous disposition. Any signal that arrives after `holdInterrupt` was called is delivered when you call `releaseInterrupt`.
- `interrupted`
  This method returns `true` if a signal caused the communicator to shut down, `false` otherwise. This allows us to distinguish intentional shutdown from a forced shutdown that was caused by a signal. This is useful, for example, for logging purposes.
- `interruptCallback`
  A subclass overrides this method to respond to signals. The method may be called concurrently with any other thread and must not raise exceptions.

By default, `Ice.Application` behaves as if `destroyOnInterrupt` was invoked, therefore our server `Main` method requires no change to ensure that the program terminates cleanly on receipt of a signal. (You can disable the signal-handling functionality of `Ice.Application` by passing the enumerator `NoSignalHandling` to the constructor. In that case, signals retain their default behavior, that is, terminate the process.) However, we add a diagnostic to report the occurrence, so our `run` method now looks like:

Copyright 2016, ZeroC, Inc.

C#

```
using System;

public class Server
{
    class App : Ice.Application
    {
        public override int run(string[] args)
        {
            // Server code here...

            if (interrupted())
                Console.Error.WriteLine(appName() + ": terminating");

            return 0;
        }
    }

    public static void Main(string[] args)
    {
        App app = new App();
        Environment.Exit(app.main(args));
    }
}
```

`Ice.Application` and Properties in C#

Apart from the functionality shown in this section, `Ice.Application` also takes care of initializing the Ice run time with property values. Properties allow you to configure the run time in various ways. For example, you can use properties to control things such as the thread pool size or port number for a server. The `main` method of `Ice.Application` is overloaded; the second version allows you to specify the name of a configuration file that will be processed during initialization.

Limitations of `Ice.Application` in C#

`Ice.Application` is a singleton class that creates a single communicator. If you are using multiple communicators, you cannot use `Ice.Application`. Instead, you must structure your code as we saw in Hello World Application (taking care to always destroy the communicator).

See Also

- Hello World Application
- Properties and Configuration
- Communicator Initialization
- Logger Facility

Copyright 2016, ZeroC, Inc.

Server-Side C-Sharp Mapping for Interfaces

The server-side mapping for interfaces provides an up-call API for the Ice run time: by implementing methods in a servant class, you provide the hook that gets the thread of control from the Ice server-side run time into your application code.

On this page:

- Skeleton Classes in C#
- Servant Classes in C#
  - Server-Side Normal and idempotent Operations in C#

Skeleton Classes in C#

On the client side, interfaces map to proxy classes. On the server side, interfaces map to *skeleton* classes. A skeleton is a class that has an abstract method for each operation on the corresponding interface. For example, consider our Slice definition for the `Node` interface:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };
    // ...
};
```

The Slice compiler generates the following definition for this interface:

Copyright 2016, ZeroC, Inc.

C#

```
namespace Filesystem
{
    public interface NodeOperations_
    {
        string name(Ice.Current __current);
    }

    public interface NodeOperationsNC_
    {
        string name();
    }

    public interface Node : Ice.Object, NodeOperations_, NodeOperationsNC_
    {
    }

    public abstract class NodeDisp_ : Ice.ObjectImpl, Node
    {
        public string name()
        {
            return name(new Ice.Current());
        }

        public abstract string name(Ice.Current __current);

        // Mapping-internal code here...
    }
}
```

The important points to note here are:

- As for the client side, Slice modules are mapped to C# namespaces with the same name, so the skeleton class definitions are part of the `Filesystem` namespace.
- For each Slice interface <*`interface-name`*>, the compiler generates C# interfaces <*`interface-name`*>`Operations_` and <*`interface-name`*>`OperationsNC_` (`NodeOperations_` and `NodeOperationsNC_` in this example). These interfaces contain a method for each operation in the Slice interface. (You can ignore the `Ice.Current` parameter for the now.)
- For each Slice interface <*`interface-name`*>, the compiler generates a C# interface <*`interface-name`*> (`Node` in this example). That interface extends `Ice.Object` and the two operations interfaces.
- For each Slice interface <*`interface-name`*>, the compiler generates an abstract class <*`interface-name`*>`Disp_` (`NodeDisp_` in this example). This abstract class is the actual skeleton class; it is the base class from which you derive your servant class.

Servant Classes in C#

In order to provide an implementation for an Ice object, you must create a servant class that inherits from the corresponding skeleton class. For example, to create a servant for the `Node` interface, you could write:

Copyright 2016, ZeroC, Inc.

C#

```
public class NodeI : NodeDisp_
{
    public NodeI(string name)
    {
        _name = name;
    }

    public override string name(Ice.Current current)
    {
        return _name;
    }

    private string _name;
}
```

By convention, servant classes have the name of their interface with an `I`-suffix, so the servant class for the `Node` interface is called `NodeI`. (This is a convention only: as far as the Ice run time is concerned, you can choose any name you prefer for your servant classes.) Note that `NodeI` extends `NodeDisp_`, that is, it derives from its skeleton class.

As far as Ice is concerned, the `NodeI` class must implement only a single method: the abstract `name` method that it inherits from its skeleton. This makes the servant class a concrete class that can be instantiated. You can add other methods and data members as you see fit to support your implementation. For example, in the preceding definition, we added a `_name` member and a constructor. (Obviously, the constructor initializes the `_name` member and the `name` method returns its value.)

Server-Side Normal and `idempotent` Operations in C#

Whether an operation is an ordinary operation or an `idempotent` operation has no influence on the way the operation is mapped. To illustrate this, consider the following interface:

Slice

```
interface Example {
   void              normalOp();
   idempotent void   idempotentOp();
};
```

The operations class for this interface looks like this:

C#

```
public interface ExampleOperations_
{
    void normalOp(Ice.Current __current);
    void idempotentOp(Ice.Current __current);
}
```


Copyright 2016, ZeroC, Inc.

Note that the signatures of the methods are unaffected by the `idempotent` qualifier.

See Also

- Slice for a Simple File System
- Parameter Passing in C-Sharp
- Raising Exceptions in C-Sharp
- Tie Classes in C-Sharp
- The Current Object

Copyright 2016, ZeroC, Inc.

Parameter Passing in C-Sharp

Parameter Passing in C#

For each parameter of a Slice operation, the C# mapping generates a corresponding parameter for the corresponding method in the *<interface-name>*`Operations_` interface. In addition, every operation has an additional, trailing parameter of type `Ice.Current`. For example, the `name` operation of the `Node` interface has no parameters, but the `name` method of the `NodeOperations_` interface has a single parameter of type `Ice.Current`. We will ignore this parameter for now.

Parameter passing on the server side follows the rules for the client side. To illustrate the rules, consider the following interface that passes string parameters in all possible directions:

Slice

```
module M {
    interface Example {
        string op(string sin, out string sout);
    };
};
```

The generated method for `op` looks as follows:

C#

```
public interface ExampleOperations_
{
    string op(string sin, out string sout, Ice.Current __current);
}
```

As you can see, there are no surprises here. For example, we could implement `op` as follows:

C#

```
using System;

public class ExampleI : ExampleDisp_
{
    public override string op(string sin, out string sout, Ice.Current current)
    {
        Console.WriteLine(sin);      // In params are initialized
        sout = "Hello World!";       // Assign out param
        return "Done";
    }
}
```

This code is in no way different from what you would normally write if you were to pass strings to and from a method; the fact that remote procedure calls are involved does not affect your code in any way. The same is true for parameters of other types, such as proxies, classes, or dictionaries: the parameter passing conventions follow normal C# rules and do not require special-purpose API calls.

Copyright 2016, ZeroC, Inc.

See Also

- Server-Side C-Sharp Mapping for Interfaces
- Raising Exceptions in C-Sharp
- Tie Classes in C-Sharp
- The Current Object

Copyright 2016, ZeroC, Inc.

Raising Exceptions in C-Sharp

To throw an exception from an operation implementation, you simply instantiate the exception, initialize it, and throw it. For example:

C#

```
// ...
public override void write(string[] text, Ice.Current current)
{
    try
    {
        // Try to write file contents here...
    }
    catch(System.Exception ex)
    {
        GenericError e = new GenericError("cannot write file", ex);
        e.reason = "Exception during write operation";
        throw e;
    }
}
```

Note that, for this example, we have supplied the optional second parameter to the `GenericError` constructor. This parameter sets the `InnerException` member of `System.Exception` and preserves the original cause of the error for later diagnosis.

If you throw an arbitrary C# run-time exception (such as an `InvalidCastException`), the Ice run time catches the exception and then returns an `UnknownException` to the client. Similarly, if you throw an "impossible" user exception (a user exception that is not listed in the exception specification of the operation), the client receives an `UnknownUserException`.

If you throw an Ice run-time exception, such `MemoryLimitException`, the client receives an `UnknownLocalException`. For that reason, you should never throw system exceptions from operation implementations. If you do, all the client will see is an `UnknownLocalException`, which does not tell the client anything useful.

Three run-time exceptions are treated specially and not changed to `UnknownLocalException` when returned to the client: `ObjectNotExistException`, `OperationNotExistException`, and `FacetNotExistException`.

See Also

- Run-Time Exceptions
- C-Sharp Mapping for Exceptions
- Server-Side C-Sharp Mapping for Interfaces
- Parameter Passing in C-Sharp
- Tie Classes in C-Sharp

Copyright 2016, ZeroC, Inc.

Tie Classes in C-Sharp

The mapping to skeleton classes requires the servant class to inherit from its skeleton class. Occasionally, this creates a problem: some class libraries require you to inherit from a base class in order to access functionality provided by the library; because C# does not support multiple implementation inheritance, this means that you cannot use such a class library to implement your servants because your servants cannot inherit from both the library class and the skeleton class simultaneously.

To allow you to still use such class libraries, Ice provides a way to write servants that replaces inheritance with delegation. This approach is supported by *tie classes*. The idea is that, instead of inheriting from the skeleton class, you simply create a class (known as an *implementation class* or *delegate class*) that contains methods corresponding to the operations of an interface. You use the `--tie` option with the `slice2cs` compiler to create a tie class. For example, the `--tie` option causes the compiler to create exactly the same code for the `Node` interface as we saw previously, but to also emit an additional tie class. For an interface `<interface-name>`, the generated tie class has the name `<interface-name>Tie_`:

Copyright 2016, ZeroC, Inc.

C#

```
public class NodeTie_ : NodeDisp_, Ice.TieBase
{
    public NodeTie_()
    {
    }

    public NodeTie_(NodeOperations_ del)
    {
        _ice_delegate = del;
    }

    public object ice_delegate()
    {
        return _ice_delegate;
    }

    public void ice_delegate(object del)
    {
        _ice_delegate = (NodeOperations_)del;
    }

    public override int GetHashCode()
    {
        return _ice_delegate == null ? 0 : _ice_delegate.GetHashCode();
    }

    public override bool Equals(object rhs)
    {
       if (object.ReferenceEquals(this, rhs))
       {
           return true;
       }
       if (!(rhs is NodeTie_))
       {
           return false;
       }
       if (_ice_delegate == null)
       {
           return ((NodeTie_)rhs)._ice_delegate == null;
       }
       return _ice_delegate.Equals(((NodeTie_)rhs)._ice_delegate);
    }

    public override string name(Ice.Current __current)
    {
        return _ice_delegate.name(__current);
    }

    private NodeOperations_ _ice_delegate;
}
```


Copyright 2016, ZeroC, Inc.

This looks a lot worse than it is: in essence, the generated tie class is simply a servant class (it extends `NodeDisp_`) that delegates each invocation of a method that corresponds to a Slice operation to your implementation class:

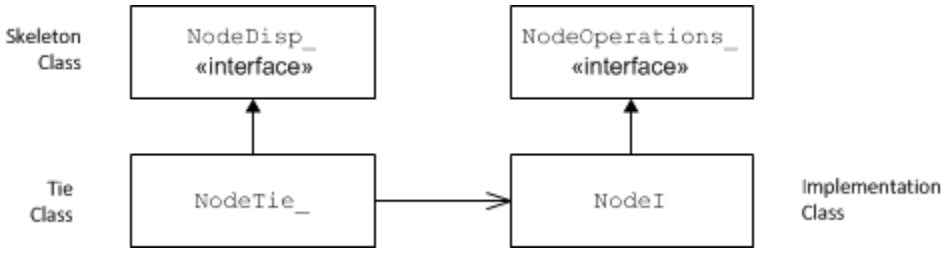

A skeleton class, tie class, and implementation class.

The `Ice.TieBase` interface defines the `ice_delegate` methods that allow you to get and set the delegate.

Given this machinery, we can create an implementation class for our `Node` interface as follows:

C#

```
public class NodeI : NodeOperations_
{
    public NodeI(string name)
    {
        _name = name;
    }

    public override string name(Ice.Current current)
    {
        return _name;
    }

    private string _name;
}
```

Note that this class is identical to our previous implementation, except that it implements the `NodeOperations_` interface and does not extend `NodeDisp_` (which means that you are now free to extend any other class to support your implementation).

To create a servant, you instantiate your implementation class and the tie class, passing a reference to the implementation instance to the tie constructor:

C#

```
NodeI fred = new NodeI("Fred");         // Create implementation
NodeTie_ servant = new NodeTie_(fred);  // Create tie
```

Alternatively, you can also default-construct the tie class and later set its delegate instance by calling `ice_delegate`:

Copyright 2016, ZeroC, Inc.

C#

```
NodeTie_ servant = new NodeTie_();      // Create tie
// ...
NodeI fred = new NodeI("Fred");         // Create implementation
// ...
servant.ice_delegate(fred);             // Set delegate
```

When using tie classes, it is important to remember that the tie instance is the servant, not your delegate. Furthermore, you must not use a tie instance to incarnate an Ice object until the tie has a delegate. Once you have set the delegate, you must not change it for the lifetime of the tie; otherwise, undefined behavior results.

You should use the tie approach only if you need to, that is, if you need to extend some base class in order to implement your servants: using the tie approach is more costly in terms of memory because each Ice object is incarnated by two C# objects (the tie and the delegate) instead of just one. In addition, call dispatch for ties is marginally slower than for ordinary servants because the tie forwards each operation to the delegate, that is, each operation invocation requires two function calls instead of one.

Also note that, unless you arrange for it, there is no way to get from the delegate back to the tie. If you need to navigate back to the tie from the delegate, you can store a reference to the tie in a member of the delegate. (The reference can, for example, be initialized by the constructor of the delegate.)

See Also

- Server-Side C-Sharp Mapping for Interfaces
- Parameter Passing in C-Sharp
- Raising Exceptions in C-Sharp
- Object Incarnation in C-Sharp

Copyright 2016, ZeroC, Inc.

Object Incarnation in C-Sharp

Having created a servant class such as the rudimentary `NodeI` class, you can instantiate the class to create a concrete servant that can receive invocations from a client. However, merely instantiating a servant class is insufficient to incarnate an object. Specifically, to provide an implementation of an Ice object, you must take the following steps:

1. Instantiate a servant class.
2. Create an identity for the Ice object incarnated by the servant.
3. Inform the Ice run time of the existence of the servant.
4. Pass a proxy for the object to a client so the client can reach it.

On this page:

- Instantiating a C# Servant
- Creating an Identity in C#
- Activating a C# Servant
- UUIDs as Identities in C#
- Creating Proxies in C#
  - Proxies and Servant Activation in C#
  - Direct Proxy Creation in C#

Instantiating a C# Servant

Instantiating a servant means to allocate an instance:

C#

```
Node servant = new NodeI("Fred");
```

This code creates a new `NodeI` instance and assigns its address to a reference of type `Node`. This works because `NodeI` is derived from `Node`, so a `Node` reference can refer to an instance of type `NodeI`. However, if we want to invoke a method of the `NodeI` class at this point, we must use a `NodeI` reference:

C#

```
NodeI servant = new NodeI("Fred");
```

Whether you use a `Node` or a `NodeI` reference depends purely on whether you want to invoke a method of the `NodeI` class: if not, a `Node` reference works just as well as a `NodeI` reference.

Creating an Identity in C#

Each Ice object requires an identity. That identity must be unique for all servants using the same object adapter.

The Ice object model assumes that all objects (regardless of their adapter) have a globally unique identity.

An Ice object identity is a structure with the following Slice definition:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    struct Identity {
        string name;
        string category;
    };
    // ...
};
```

The full identity of an object is the combination of both the `name` and `category` fields of the `Identity` structure. For now, we will leave the `category` field as the empty string and simply use the `name` field. (The `category` field is most often used in conjunction with servant locators.)

To create an identity, we simply assign a key that identifies the servant to the `name` field of the `Identity` structure:

C#

```
Ice.Identity id = new Ice.Identity();
id.name = "Fred"; // Not unique, but good enough for now
```

Activating a C# Servant

Merely creating a servant instance does nothing: the Ice run time becomes aware of the existence of a servant only once you explicitly tell the object adapter about the servant. To activate a servant, you invoke the `add` operation on the object adapter. Assuming that we have access to the object adapter in the `_adapter` variable, we can write:

C#

```
_adapter.add(servant, id);
```

Note the two arguments to `add`: the servant and the object identity. Calling `add` on the object adapter adds the servant and the servant's identity to the adapter's servant map and links the proxy for an Ice object to the correct servant instance in the server's memory as follows:

1. The proxy for an Ice object, apart from addressing information, contains the identity of the Ice object. When a client invokes an operation, the object identity is sent with the request to the server.
2. The object adapter receives the request, retrieves the identity, and uses the identity as an index into the servant map.
3. If a servant with that identity is active, the object adapter retrieves the servant from the servant map and dispatches the incoming request into the correct method on the servant.

Assuming that the object adapter is in the active state, client requests are dispatched to the servant as soon as you call `add`.

UUIDs as Identities in C#

The Ice object model assumes that object identities are globally unique. One way of ensuring that uniqueness is to use UUIDs (Universally Unique Identifiers) as identities. .NET provides a helper function that we can use to create such identities:

Copyright 2016, ZeroC, Inc.

C#

```
public class Example
{
    public static void Main(string[] args)
    {
        System.Console.WriteLine(System.Guid.NewGuid().ToString());
    }
}
```

When executed, this program prints a unique string such as `5029a22c-e333-4f87-86b1-cd5e0fcce509`. Each call to `NewGuid` creates a string that differs from all previous ones.

You can use a UUID such as this to create object identities. For convenience, the object adapter has an operation `addWithUUID` that generates a UUID and adds a servant to the servant map in a single step. Using this operation, we can create an identity and register a servant with that identity in a single step as follows:

C#

```
_adapter.addWithUUID(new NodeI("Fred"));
```

Creating Proxies in C#

Once we have activated a servant for an Ice object, the server can process incoming client requests for that object. However, clients can only access the object once they hold a proxy for the object. If a client knows the server's address details and the object identity, it can create a proxy from a string, as we saw in our first example in Hello World Application. However, creation of proxies by the client in this manner is usually only done to allow the client access to initial objects for bootstrapping. Once the client has an initial proxy, it typically obtains further proxies by invoking operations.

The object adapter contains all the details that make up the information in a proxy: the addressing and protocol information, and the object identity. The Ice run time offers a number of ways to create proxies. Once created, you can pass a proxy to the client as the return value or as an out-parameter of an operation invocation.

Proxies and Servant Activation in C#

The `add` and `addWithUUID` servant activation operations on the object adapter return a proxy for the corresponding Ice object. This means we can write:

C#

```
NodePrx proxy = NodePrxHelper.uncheckedCast(_adapter.addWithUUID(new NodeI("Fred"))
);
```

Here, `addWithUUID` both activates the servant and returns a proxy for the Ice object incarnated by that servant in a single step.

Note that we need to use an `uncheckedCast` here because `addWithUUID` returns a proxy of type `Ice.ObjectPrx`.

Direct Proxy Creation in C#

The object adapter offers an operation to create a proxy for a given identity:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local interface ObjectAdapter {
        Object* createProxy(Identity id);
        // ...
    };
};
```

Note that `createProxy` creates a proxy for a given identity whether a servant is activated with that identity or not. In other words, proxies have a life cycle that is quite independent from the life cycle of servants:

C#

```
Ice.Identity id = new Ice.Identity();
id.name = System.Guid.NewGuid().ToString();
Ice.ObjectPrx o = _adapter.createProxy(id);
```

This creates a proxy for an Ice object with the identity returned by `NewGuid`. Obviously, no servant yet exists for that object so, if we return the proxy to a client and the client invokes an operation on the proxy, the client will receive an `ObjectNotExistException`. (We examine these life cycle issues in more detail in Object Life Cycle.)

See Also

- Hello World Application
- Server-Side C-Sharp Mapping for Interfaces
- Object Adapter States
- Servant Locators
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

Asynchronous Method Dispatch (AMD) in C-Sharp

The number of simultaneous synchronous requests a server is capable of supporting is determined by the number of threads in the server's thread pool. If all of the threads are busy dispatching long-running operations, then no threads are available to process new requests and therefore clients may experience an unacceptable lack of responsiveness.

Asynchronous Method Dispatch (AMD), the server-side equivalent of AMI, addresses this scalability issue. Using AMD, a server can receive a request but then suspend its processing in order to release the dispatch thread as soon as possible. When processing resumes and the results are available, the server sends a response explicitly using a callback object provided by the Ice run time.

AMD is transparent to the client, that is, there is no way for a client to distinguish a request that, in the server, is processed synchronously from a request that is processed asynchronously.

In practical terms, an AMD operation typically queues the request data (i.e., the callback object and operation arguments) for later processing by an application thread (or thread pool). In this way, the server minimizes the use of dispatch threads and becomes capable of efficiently supporting thousands of simultaneous clients.

An alternate use case for AMD is an operation that requires further processing after completing the client's request. In order to minimize the client's delay, the operation returns the results while still in the dispatch thread, and then continues using the dispatch thread for additional work.

On this page:

- Enabling AMD with Metadata in C#
- AMD Mapping in C#
  - Callback Interface for AMD
  - Dispatch Method for AMD
- AMD Exceptions in C#
- AMD Example in C#

Enabling AMD with Metadata in C#

To enable asynchronous dispatch, you must add an `["amd"]` metadata directive to your Slice definitions. The directive applies at the interface and the operation level. If you specify `["amd"]` at the interface level, all operations in that interface use asynchronous dispatch; if you specify `["amd"]` for an individual operation, only that operation uses asynchronous dispatch. In either case, the metadata directive replaces synchronous dispatch, that is, a particular operation implementation must use synchronous or asynchronous dispatch and cannot use both.

Consider the following Slice definitions:

Slice

```
["amd"] interface I {
  bool isValid();
  float computeRate();
};

interface J {
  ["amd"] void startProcess();
  int endProcess();
};
```

In this example, both operations of interface `I` use asynchronous dispatch, whereas, for interface `J`, `startProcess` uses asynchronous dispatch and `endProcess` uses synchronous dispatch.

Specifying metadata at the operation level (rather than at the interface or class level) minimizes the amount of generated code and, more importantly, minimizes complexity: although the asynchronous model is more flexible, it is also more complicated to use. It is therefore in your best interest to limit the use of the asynchronous model to those operations that need it, while using the simpler synchronous model for the rest.

AMD Mapping in C#

The C# mapping emits the following code for each AMD operation:

Copyright 2016, ZeroC, Inc.

1. Callback interface
2. Dispatch method

Callback Interface for AMD

A callback interface is used by the implementation to notify the Ice run time about the completion of an operation. The name of this interface is formed using the pattern `AMD_`*`class_op`*. For example, an operation named `foo` defined in interface `I` results in an interface named `AMD_I_foo`. The interface is generated in the same scope as the interface or class containing the operation. Two methods are provided:

C#

```
public void ice_response(<params>);
```

The `ice_response` method allows the server to report the successful completion of the operation. If the operation has a non-`void` return type, the first parameter to `ice_response` is the return value. Parameters corresponding to the operation's `out` parameters follow the return value, in the order of declaration.

C#

```
public void ice_exception(System.Exception ex);
```

The `ice_exception` method allows the server to raise an exception. Although any exception type could conceivably be passed to `ice_exception`, the Ice run time validates the exception value using the same semantics as for synchronous dispatch.

Neither `ice_response` nor `ice_exception` throw any exceptions to the caller.

Dispatch Method for AMD

The dispatch method, whose name has the suffix `_async`, has a `void` return type. The first parameter is a reference to an instance of the callback interface described above. The remaining parameters comprise the `in` parameters of the operation, in the order of declaration.

For example, suppose we have defined the following operation:

Slice

```
interface I {
  ["amd"] int foo(short s, out long l);
};
```

The callback interface generated for operation `foo` is shown below:

Copyright 2016, ZeroC, Inc.

C#

```
public interface AMD_I_foo
{
    void ice_response(int __ret, long l);
    void ice_exception(System.Exception ex);
}
```

The dispatch method for asynchronous invocation of operation `foo` is generated as follows:

C#

```
public abstract void foo_async(AMD_I_foo __cb, short s, Ice.Current __current);
```

AMD Exceptions in C#

There are two processing contexts in which the logical implementation of an AMD operation may need to report an exception: the dispatch thread (the thread that receives the invocation), and the response thread (the thread that sends the response).

These are not necessarily two different threads: it is legal to send the response from the dispatch thread.

Although we recommend that the callback object be used to report all exceptions to the client, it is legal for the implementation to raise an exception instead, but only from the dispatch thread.

As you would expect, an exception raised from a response thread cannot be caught by the Ice run time; the application's run time environment determines how such an exception is handled. Therefore, a response thread must ensure that it traps all exceptions and sends the appropriate response using the callback object. Otherwise, if a response thread is terminated by an uncaught exception, the request may never be completed and the client might wait indefinitely for a response.

Whether raised in a dispatch thread or reported via the callback object, user exceptions are validated as described in User Exceptions, and local exceptions may undergo the translation described in Run-Time Exceptions.

AMD Example in C#

To demonstrate the use of AMD in Ice, let us define the Slice interface for a simple computational engine:

Copyright 2016, ZeroC, Inc.

Slice

```
module Demo {
    sequence<float> Row;
    sequence<Row> Grid;

    exception RangeError {};

    interface Model {
        ["amd"] Grid interpolate(Grid data, float factor)
            throws RangeError;
    };
};
```

Given a two-dimensional grid of floating point values and a factor, the `interpolate` operation returns a new grid of the same size with the values interpolated in some interesting (but unspecified) way.

Our servant class derives from `Demo._ModelDisp` and supplies a definition for the `interpolate_async` method that creates a `Job` to hold the callback object and arguments, and adds the `Job` to a queue. The method uses a `lock` statement to guard access to the queue:

C#

```
public class ModelI : Demo.ModelDisp_
{
    public override void interpolate_async(
        Demo.AMD_Model_interpolate cb,
        float[][] data,
        float factor,
        Ice.Current current)
    {
        lock(this)
        {
            _jobs.Add(new Job(cb, data, factor));
        }
    }

    private System.Collections.ArrayList _jobs
= new System.Collections.ArrayList();
}
```

After queuing the information, the operation returns control to the Ice run time, making the dispatch thread available to process another request. An application thread removes the next `Job` from the queue and invokes `execute`, which uses `interpolateGrid` (not shown) to perform the computational work:

Copyright 2016, ZeroC, Inc.

C#

```
public class Job {
    public Job(Demo.AMD_Model_interpolate cb, float[][] grid, float factor)
    {
        _cb = cb;
        _grid = grid;
        _factor = factor;
    }

    public void execute()
    {
        if (!interpolateGrid()) {
            _cb.ice_exception(new Demo.RangeError());
            return;
        }
        _cb.ice_response(_grid);
    }

    private boolean interpolateGrid()
    {
        // ...
    }

    private Demo.AMD_Model_interpolate _cb;
    private float[][] _grid;
    private float _factor;
}
```

If `interpolateGrid` returns `false`, then `ice_exception` is invoked to indicate that a range error has occurred. The `return` statement following the call to `ice_exception` is necessary because `ice_exception` does not throw an exception; it only marshals the exception argument and sends it to the client.

If interpolation was successful, `ice_response` is called to send the modified grid back to the client.

See Also

- User Exceptions
- Run-Time Exceptions
- Asynchronous Method Invocation (AMI) in C-Sharp
- The Ice Threading Model

Copyright 2016, ZeroC, Inc.

Example of a File System Server in C-Sharp

This page presents the source code for a C# server that implements our file system and communicates with the client we wrote earlier. The code is fully functional, apart from the required interlocking for threads.

The server is free of code that relates to distribution: most of the server code is simply application logic that would be present just the same for a non-distributed version. Again, this is one of the major advantages of Ice: distribution concerns are kept away from application code so that you can concentrate on developing application logic instead of networking infrastructure.

The server code shown here is not quite correct as it stands: if two clients access the same file in parallel, each via a different thread, one thread may read the `_lines` data member while another thread updates it. Obviously, if that happens, we may write or return garbage or, worse, crash the server. However, it is trivial to make the `read` and `write` operations thread-safe. We discuss thread safety in The Ice Threading Model.

On this page:

- Implementing a File System Server in C#
- Server Main Program in C#
- FileI Servant Class in C#
- DirectoryI Servant Class in C#
  - DirectoryI Data Members
  - DirectoryI Constructor
  - DirectoryI Methods

Implementing a File System Server in C#

We have now seen enough of the server-side C# mapping to implement a server for our file system. (You may find it useful to review these Slice definitions before studying the source code.)

Our server is composed of three source files:

- `Server.cs`
  This file contains the server main program.
- `DirectoryI.cs`
  This file contains the implementation for the `Directory` servants.
- `FileI.cs`
  This file contains the implementation for the `File` servants.

Server Main Program in C#

Our server main program, in the file `Server.cs`, uses the `Ice.Application` class. The `run` method installs a signal handler, creates an object adapter, instantiates a few servants for the directories and files in the file system, and then activates the adapter. This leads to a main program as follows:

C#

```
using Filesystem;
using System;

public class Server
{
    class App : Ice.Application
    {
        public override int run(string[] args)
        {
            // Terminate cleanly on receipt of a signal
            //
            shutdownOnInterrupt();
```


Copyright 2016, ZeroC, Inc.

```
            // Create an object adapter (stored in the _adapter
            // static members)
            //
            Ice.ObjectAdapter adapter
= communicator().createObjectAdapterWithEndpoints(
                "SimpleFilesystem", "default -p 10000");
            DirectoryI._adapter = adapter;
            FileI._adapter = adapter;

            // Create the root directory (with name "/" and no
            // parent)
            //
            DirectoryI root = new DirectoryI("/", null);

            // Create a file called "README" in the root directory
            //
            File file = new FileI("README", root);
            string[] text;
            text = new string[] {
                "This file system contains a collection of poetry."
            };
            try {
                file.write(text);
            } catch (GenericError e) {
                Console.Error.WriteLine(e.reason);
            }

            // Create a directory called "Coleridge"
            // in the root directory
            //
            DirectoryI coleridge = new DirectoryI("Coleridge", root);

            // Create a file called "Kubla_Khan"
            // in the Coleridge directory
            //
            file = new FileI("Kubla_Khan", coleridge);
            text = new string[] { "In Xanadu did Kubla Khan",
                                  "A stately pleasure-dome decree:",
                                  "Where Alph, the sacred river, ran",
                                  "Through caverns measureless to man",
                                  "Down to a sunless sea." };
            try {
                file.write(text);
            } catch (GenericError e) {
                Console.Error.WriteLine(e.reason);
            }

            // All objects are created, allow client requests now
            //
            adapter.activate();

            // Wait until we are done
            //
            communicator().waitForShutdown();

            if (interrupted())
```


Copyright 2016, ZeroC, Inc.

```
            Console.Error.WriteLine(appName() + ": terminating");

        return 0;
    }
}

public static void Main(string[] args)
{
    App app = new App();
    Environment.Exit(app.main(args));
```


Copyright 2016, ZeroC, Inc.

```
    }
}
```

The code uses a `using` directive for the `Filesystem` namespace. This avoids having to continuously use fully-qualified identifiers with a `Filesystem.` prefix.

The next part of the source code is the definition of the `Server` class, which includes a nested class that derives from `Ice.Application` and contains the main application logic in its `run` method. Much of this code is boiler plate that we saw previously: we create an object adapter, and, towards the end, activate the object adapter and call `waitForShutdown`.

The interesting part of the code follows the adapter creation: here, the server instantiates a few nodes for our file system to create the structure shown below:

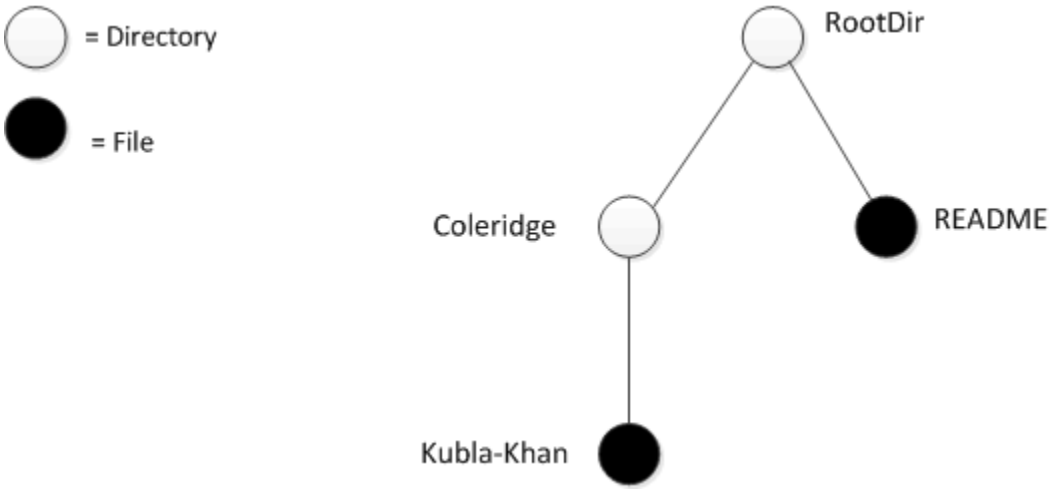

A small file system.

As we will see shortly, the servants for our directories and files are of type `DirectoryI` and `FileI`, respectively. The constructor for either type of servant accepts two parameters, the name of the directory or file to be created and a reference to the servant for the parent directory. (For the root directory, which has no parent, we pass a null parent.) Thus, the statement

C#

```
DirectoryI root = new DirectoryI("/", null);
```

creates the root directory, with the name `"/"` and no parent directory.

Here is the code that establishes the structure in the above illustration:

Copyright 2016, ZeroC, Inc.

C#

```
// Create the root directory (with name "/" and no parent)
//
DirectoryI root = new DirectoryI("/", null);

// Create a file called "README" in the root directory
//
File file = new FileI("README", root);
string[] text;
text = new string[] {
    "This file system contains a collection of poetry."
};
try {
    file.write(text);
} catch (GenericError e) {
    Console.Error.WriteLine(e.reason);
}

// Create a directory called "Coleridge"
// in the root directory
//
DirectoryI coleridge = new DirectoryI("Coleridge", root);

// Create a file called "Kubla_Khan"
// in the Coleridge directory
//
file = new FileI("Kubla_Khan", coleridge);
text = new string[] { "In Xanadu did Kubla Khan",
                      "A stately pleasure-dome decree:",
                      "Where Alph, the sacred river, ran",
                      "Through caverns measureless to man",
                      "Down to a sunless sea." };
try {
    file.write(text);
} catch (GenericError e) {
    Console.Error.WriteLine(e.reason);
}
```

We first create the root directory and a file `README` within the root directory. (Note that we pass a reference to the root directory as the parent when we create the new node of type `FileI`.)

The next step is to fill the file with text:

Copyright 2016, ZeroC, Inc.

C#

```
        string[] text;
        text = new string[] {
            "This file system contains a collection of poetry."
        };
        try {
            file.write(text);
        } catch (GenericError e) {
            Console.Error.WriteLine(e.reason);
        }
```

Recall that Slice sequences by default map to C# arrays. The Slice type `Lines` is simply an array of strings; we add a line of text to our `README` file by initializing the `text` array to contain one element.

Finally, we call the Slice `write` operation on our `FileI` servant by writing:

C#

```
            file.write(text);
```

This statement is interesting: the server code invokes an operation on one of its own servants. Because the call happens via a reference to the servant (of type `FileI`) and not via a proxy (of type `FilePrx`), the Ice run time does not know that this call is even taking place — such a direct call into a servant is not mediated by the Ice run time in any way and is dispatched as an ordinary C# method call.

In similar fashion, the remainder of the code creates a subdirectory called `Coleridge` and, within that directory, a file called `Kubla_Khan` to complete the structure in the above illustration.

`FileI` Servant Class in C#

Our `FileI` servant class has the following basic structure:

C#

```
using Filesystem;
using System;

public class FileI : FileDisp_
{
    // Constructor and operations here...

    public static Ice.ObjectAdapter _adapter;
    private string _name;
    private DirectoryI _parent;
    private string[] _lines;
}
```

The class has a number of data members:

Copyright 2016, ZeroC, Inc.

- `_adapter`
  This static member stores a reference to the single object adapter we use in our server.
- `_name`
  This member stores the name of the file incarnated by the servant.
- `_parent`
  This member stores the reference to the servant for the file's parent directory.
- `_lines`
  This member holds the contents of the file.

The `_name` and `_parent` data members are initialized by the constructor:

C#

```
public FileI(string name, DirectoryI parent)
{
    _name = name;
    _parent = parent;

    Debug.Assert(_parent != null);

    // Create an identity
    //
    Ice.Identity myID = new Ice.Identity();
    myId.name = System.Guid.NewGuid().ToString();

    // Add the identity to the object adapter
    //
    _adapter.add(this, myID);

    // Create a proxy for the new node and
    // add it as a child to the parent
    //
    NodePrx thisNode = NodePrxHelper.uncheckedCast(_adapter.createProxy(myID));
    _parent.addChild(thisNode);
}
```

After initializing the `_name` and `_parent` members, the code verifies that the reference to the parent is not null because every file must have a parent directory. The constructor then generates an identity for the file by calling `NewGuid` and adds itself to the servant map by calling `ObjectAdapter.add`. Finally, the constructor creates a proxy for this file and calls the `addChild` method on its parent directory. `addChild` is a helper function that a child directory or file calls to add itself to the list of descendant nodes of its parent directory. We will see the implementation of this function in `DirectoryI` Methods.

The remaining methods of the `FileI` class implement the Slice operations we defined in the `Node` and `File` Slice interfaces:

Copyright 2016, ZeroC, Inc.

C#

```
    // Slice Node::name() operation

    public override string name(Ice.Current current)
    {
        return _name;
    }

    // Slice File::read() operation

    public override string[] read(Ice.Current current)
    {
        return _lines;
    }

    // Slice File::write() operation

    public override void write(string[] text, Ice.Current current)
    {
        _lines = text;
    }
```

The `name` method is inherited from the generated `Node` interface (which is a base interface of the `_FileDisp` class from which `FileI` is derived). It simply returns the value of the `_name` member.

The `read` and `write` methods are inherited from the generated `File` interface (which is a base interface of the `_FileDisp` class from which `FileI` is derived) and simply return and set the `_lines` member.

`DirectoryI` Servant Class in C#

The `DirectoryI` class has the following basic structure:

C#

```
using Filesystem;
using System;
using System.Collections;

public class DirectoryI : DirectoryDisp_
{
    // Constructor and operations here...

    public static Ice.ObjectAdapter _adapter;
    private string _name;
    private DirectoryI _parent;
    private ArrayList _contents = new ArrayList();
}
```

`DirectoryI` Data Members

Copyright 2016, ZeroC, Inc.

As for the `FileI` class, we have data members to store the object adapter, the name, and the parent directory. (For the root directory, the `_parent` member holds a null reference.) In addition, we have a `_contents` data member that stores the list of child directories. These data members are initialized by the constructor:

C#

```
public DirectoryI(string name, DirectoryI parent)
{
    _name = name;
    _parent = parent;

    // Create an identity. The
    // parent has the fixed identity "RootDir"
    //
    Ice.Identity myID = new Ice.Identity();
    myID.name = _parent != null ? System.Guid.NewGuid().ToString() : "RootDir";

    // Add the identity to the object adapter
    //
    _adapter.add(this, myID);

    // Create a proxy for the new node and
    // add it as a child to the parent
    //
    NodePrx thisNode = NodePrxHelper.uncheckedCast(_adapter.createProxy(myID));
    if (_parent != null)
        _parent.addChild(thisNode);
}
```

`DirectoryI` Constructor

The constructor creates an identity for the new directory by calling `NewGuid`. (For the root directory, we use the fixed identity `"RootDir"`.) The servant adds itself to the servant map by calling `ObjectAdapter.add` and then creates a proxy to itself and passes it to the `addChild` helper function.

`DirectoryI` Methods

`addChild` simply adds the passed reference to the `_contents` list:

C#

```
public void addChild(NodePrx child)
{
    _contents.Add(child);
}
```

The remainder of the operations, `name` and `list`, are trivial:

Copyright 2016, ZeroC, Inc.

C#

```
public override string name(Ice.Current current)
{
    return _name;
}

public override NodePrx[] list(Ice.Current current)
{
    return (NodePrx[])_contents.ToArray(typeof(NodePrx));
}
```

Note that the `_contents` member is of type `System.Collections.ArrayList`, which is convenient for the implementation of the `addChild` method. However, this requires us to convert the list into a C# array in order to return it from the `list` operation.

See Also

- Slice for a Simple File System
- Example of a File System Client in C-Sharp
- The Server-Side main Method in C-Sharp
- C-Sharp Mapping for Sequences
- The Ice Threading Model

Copyright 2016, ZeroC, Inc.

The .NET Utility Library

Ice for .Net includes a number of utility APIs in the `Ice.Util` class. This appendix summarizes the contents of these APIs for your reference.

On this page:

- Communicator Initialization Methods
- Identity Conversion
- Per-Process Logger Methods
- Property Creation Methods
- Proxy Comparison Methods
- Stream Creation
- Version Information

Communicator Initialization Methods

`Ice.Util` provides a number of overloaded `initialize` methods that create a communicator.

Identity Conversion

`Ice.Util` contains two methods to convert object identities of type `Ice.Identity` to and from strings.

Per-Process Logger Methods

`Ice.Util` provides methods for getting and setting the per-process logger.

Property Creation Methods

`Ice.Util` provides a number of overloaded `createProperties` methods that create property sets.

Proxy Comparison Methods

Two methods, `proxyIdentityCompare` and `proxyIdentityAndFacetCompare`, allow you to compare object identities that are stored in proxies (either ignoring the facet or taking the facet into account).

Stream Creation

The methods `createInputStream`, `wrapInputStream` and `createOutputStream` create streams for use with dynamic invocation.

Version Information

The `stringVersion` and `intVersion` methods return the version of the Ice run time:

C#

```
public static string stringVersion();
public static int intVersion();
```

The `stringVersion` method returns the Ice version in the form *<major>.<minor>.<patch>*, for example, `3.4.2`. For beta releases, the version is *<major>.<minor>*b, for example, `3.4b`.

The `intVersion` method returns the Ice version in the form *AABBCC*, where *AA* is the major version number, *BB* is the minor version number, and *CC* is patch level, for example, `30402` for version 3.4.2. For beta releases, the patch level is set to 51 so, for example, for version 3.4b, the value is `30451`.

See Also

- Command-Line Parsing and Initialization

Copyright 2016, ZeroC, Inc.

- Setting Properties
- C-Sharp Streaming Interfaces
- Object Identity

Copyright 2016, ZeroC, Inc.

Java Mapping

Topics

- Client-Side Slice-to-Java Mapping
- Server-Side Slice-to-Java Mapping
- The Java Utility Library
- Custom Class Loaders
- Java Interrupts

Copyright 2016, ZeroC, Inc.

Client-Side Slice-to-Java Mapping

In this section, we present the client-side Slice-to-Java mapping. The client-side Slice-to-Java mapping defines how Slice data types are translated to Java types, and how clients invoke operations, pass parameters, and handle errors. Much of the Java mapping is intuitive. For example, Slice sequences map to Java arrays, so there is essentially nothing new you have to learn in order to use Slice sequences in Java.

The Java API to the Ice run time is fully thread-safe. Obviously, you must still synchronize access to data from different threads. For example, if you have two threads sharing a sequence, you cannot safely have one thread insert into the sequence while another thread is iterating over the sequence. However, you only need to concern yourself with concurrent access to your own data — the Ice run time itself is fully thread safe, and none of the Ice API calls require you to acquire or release a lock before you safely can make the call.

Much of what appears in this chapter is reference material. We suggest that you skim the material on the initial reading and refer back to specific sections as needed. However, we recommend that you read at least the mappings for exceptions, interfaces, and operations in detail because these sections cover how to call operations from a client, pass parameters, and handle exceptions.

In order to use the Java mapping, you should need no more than the Slice definition of your application and knowledge of the Java mapping rules. In particular, looking through the generated code in order to discern how to use the Java mapping is likely to be inefficient, due to the amount of detail. Of course, occasionally, you may want to refer to the generated code to confirm a detail of the mapping, but we recommend that you otherwise use the material presented here to see how to write your client-side code.

The `Ice` Package
All of the APIs for the Ice run time are nested in the `Ice` package, to avoid clashes with definitions for other libraries or applications. Some of the contents of the `Ice` package are generated from Slice definitions; other parts of the `Ice` package provide special-purpose definitions that do not have a corresponding Slice definition. We will incrementally cover the contents of the `Ice` package throughout the remainder of the manual.

Topics

- Java Mapping for Identifiers
- Java Mapping for Modules
- Java Mapping for Built-In Types
- Java Mapping for Enumerations
- Java Mapping for Structures
- Java Mapping for Sequences
- Java Mapping for Dictionaries
- Java Mapping for Constants
- Java Mapping for Exceptions
- Java Mapping for Interfaces
- Java Mapping for Operations
- Java Mapping for Classes
- Java Mapping for Optional Data Members
- Serializable Objects in Java
- Customizing the Java Mapping
- Asynchronous Method Invocation (AMI) in Java
- Using the Slice Compiler for Java
- slice2java Command-Line Options
- Using Slice Checksums in Java
- Example of a File System Client in Java

Copyright 2016, ZeroC, Inc.

Java Mapping for Identifiers

A Slice identifier maps to an identical Java identifier. For example, the Slice identifier `Clock` becomes the Java identifier `Clock`. There is one exception to this rule: if a Slice identifier is the same as a Java keyword or is an identifier reserved by the Ice run time (such as `checkedCast`), the corresponding Java identifier is prefixed with an underscore. For example, the Slice identifier `while` is mapped as `_while`.

You should try to avoid such identifiers as much as possible.

A single Slice identifier often results in several Java identifiers. For example, for a Slice interface named `Foo`, the generated Java code uses the identifiers `Foo` and `FooPrx` (among others). If the interface has the name `while`, the generated identifiers are `_while` and `whilePrx` (*not* `_whilePrx`), that is, the underscore prefix is applied only to those generated identifiers that actually require it.

See Also

- Lexical Rules
- Java Mapping for Modules
- Java Mapping for Built-In Types
- Java Mapping for Enumerations
- Java Mapping for Structures
- Java Mapping for Sequences
- Java Mapping for Dictionaries
- Java Mapping for Constants
- Java Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

Java Mapping for Modules

A Slice module maps to a Java package with the same name as the Slice module. The mapping preserves the nesting of the Slice definitions. For example:

Slice

```
// Definitions at global scope here...

module M1 {
    // Definitions for M1 here...
    module M2 {
        // Definitions for M2 here...
    };
};

// ...

module M1 {     // Reopen M1
    // More definitions for M1 here...
};
```

This definition maps to the corresponding Java definitions:

Java

```
package M1;
// Definitions for M1 here...

package M1.M2;
// Definitions for M2 here...

package M1;
// Definitions for M1 here...
```

Note that these definitions appear in the appropriate source files; source files for definitions in module `M1` are generated in directory `M1` under neath the top-level directory, and source files for definitions for module `M2` are generated in directory `M1/M2` underneath the top-level directory. You can set the top-level output directory using the `--output-dir` option with `slice2java`.

See Also

- Modules
- Using the Slice Compilers
- Java Mapping for Identifiers
- Java Mapping for Built-In Types
- Java Mapping for Enumerations
- Java Mapping for Structures
- Java Mapping for Sequences
- Java Mapping for Dictionaries
- Java Mapping for Constants
- Java Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

Java Mapping for Built-In Types

The Slice built-in types are mapped to Java types as follows:

Slice	Java
bool	`boolean`
byte	`byte`
short	`short`
int	`int`
long	`long`
float	`float`
double	`double`
string	`String`

Mapping of Slice built-in types to Java.

See Also

- Basic Types
- Java Mapping for Identifiers
- Java Mapping for Modules
- Java Mapping for Enumerations
- Java Mapping for Structures
- Java Mapping for Sequences
- Java Mapping for Dictionaries
- Java Mapping for Constants
- Java Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

Java Mapping for Enumerations

A Slice enumeration maps to the corresponding enumeration in Java. For example:

Slice

```
enum Fruit { Apple, Pear, Orange };
```

The Java mapping for `Fruit` is shown below:

Java

```
public enum Fruit implements java.io.Serializable {
    Apple,
    Pear,
    Orange;

    public int value();

    public static Fruit valueOf(int v);

    // ...
}
```

Given the above definitions, we can use enumerated values as follows:

Java

```
Fruit f1 = Fruit.Apple;
Fruit f2 = Fruit.Orange;

if (f1 == Fruit.Apple) // Compare with constant
    // ...

if (f1 == f2)          // Compare two enums
    // ...

switch (f2) {          // Switch on enum
case Fruit.Apple:
    // ...
    break;
case Fruit.Pear:
    // ...
    break;
case Fruit.Orange:
    // ...
    break;
}
```

The Java mapping includes two methods of interest. The `value` method returns the Slice value of an enumerator, which is not necessarily the same as its ordinal value. The `valueOf` method translates a Slice value into its corresponding enumerator, or returns `null` if no match

Copyright 2016, ZeroC, Inc.

is found. Note that the generated class contains a number of other members, which we have not shown. These members are internal to the Ice run time and you must not use them in your application code (because they may change from release to release).

In the `Fruit` definition above, the Slice value of each enumerator matches its ordinal value. This will not be true if we modify the definition to include a custom enumerator value:

Slice

```
enum Fruit { Apple, Pear = 3, Orange };
```

The table below shows the new relationship between ordinal value and Slice value:

Enumerator	Ordinal	Slice
`Apple`	0	0
`Pear`	1	3
`Orange`	2	4

Java enumerated types inherit implicitly from `java.lang.Enum`, which defines methods such as `ordinal` and `compareTo` that operate on the *ordinal* value of an enumerator, not its Slice value.

See Also

- Enumerations
- Java Mapping for Structures
- Java Mapping for Sequences
- Java Mapping for Dictionaries

Copyright 2016, ZeroC, Inc.

Java Mapping for Structures

On this page:

- Basic Java Mapping for Structures
- Java Default Constructors for Structures

Basic Java Mapping for Structures

A Slice structure maps to a Java class with the same name. For each Slice data member, the Java class contains a corresponding public data member. For example, here is our Employee structure once more:

Slice

```
struct Employee {
    long number;
    string firstName;
    string lastName;
};
```

The Slice-to-Java compiler generates the following definition for this structure:

Copyright 2016, ZeroC, Inc.

Java

```
public final class Employee implements java.lang.Cloneable, java.io.Serializable {
    public long number;
    public String firstName;
    public String lastName;

    public Employee() {}

    public Employee(long number, String firstName, String lastName) {
        this.number = number;
        this.firstName = firstName;
        this.lastName = lastName;
    }

    public boolean equals(java.lang.Object rhs) {
        // ...
    }
    public int hashCode() {
        // ...
    }

    public java.lang.Object clone()
        java.lang.Object o;
        try
        {
            o = super.clone();
        }
        catch(java.lang.CloneNotSupportedException ex)
        {
            assert false; // impossible
        }
        return o;
    }
}
```

For each data member in the Slice definition, the Java class contains a corresponding public data member of the same name. Note that you can optionally customize the mapping for data members to use getters and setters instead.

The `equals` member function compares two structures for equality. Note that the generated class also provides the usual `hashCode` and `clone` methods. (`clone` has the default behavior of making a shallow copy.)

Java Default Constructors for Structures

Structures have a default constructor that initializes data members as follows:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero

Copyright 2016, ZeroC, Inc.

`bool`	False
`sequence`	Null
`dictionary`	Null
`class/interface`	Null

The constructor won't explicitly initialize a data member if the default Java behavior for that type produces the desired results.

If you wish to ensure that data members of primitive and enumerated types are initialized to specific values, you can declare default values in your Slice definition. The default constructor initializes each of these data members to its declared value instead.

Structures also have a second constructor that has one parameter for each data member. This allows you to construct and initialize an instance in a single statement (instead of first having to construct the instance and then assign to its members).

See Also

- Structures
- Java Mapping for Enumerations
- Java Mapping for Sequences
- Java Mapping for Dictionaries
- Customizing the Java Mapping

Copyright 2016, ZeroC, Inc.

Java Mapping for Sequences

A Slice sequence maps to a Java array. This means that the Slice-to-Java compiler does not generate a separate named type for a Slice sequence.

For example:

Slice

```
sequence<Fruit> FruitPlatter;
```

This definition simply corresponds to the Java type `Fruit[]`. Naturally, because Slice sequences are mapped to Java arrays, you can take advantage of all the array functionality provided by Java, such as initialization, assignment, cloning, and the `length` member. For example:

Java

```
Fruit[] platter = { Fruit.Apple, Fruit.Pear };
assert(platter.length == 2);
```

Alternate mappings for sequence types are also possible.

See Also

- Sequences
- Java Mapping for Enumerations
- Java Mapping for Structures
- Java Mapping for Dictionaries
- Customizing the Java Mapping

Copyright 2016, ZeroC, Inc.

Java Mapping for Dictionaries

Here is the definition of our EmployeeMap once more:

Slice

```
dictionary<long, Employee> EmployeeMap;
```

As for sequences, the Java mapping does not create a separate named type for this definition. Instead, the dictionary is simply an instance of the generic type `java.util.Map<K, V>`, where *K* is the mapping of the key type and *V* is the mapping of the value type. In the example above, `EmployeeMap` is mapped to the Java type `java.util.Map<Long, Employee>`. The following code demonstrates how to allocate and use an instance of `EmployeeMap`:

Java

```
java.util.Map<Long, Employee> em = new java.util.HashMap<Long, Employee>();

Employee e = new Employee();
e.number = 31;
e.firstName = "James";
e.lastName = "Gosling";

em.put(e.number, e);
```

The type-safe nature of the mapping makes iterating over the dictionary quite convenient:

Java

```
for (java.util.Map.Entry<Long, Employee> i : em.entrySet()) {
    long num = i.getKey();
    Employee emp = i.getValue();
    System.out.println(emp.firstName + " was employee #" + num);
}
```

Alternate mappings for dictionary types are also possible.

See Also

- Dictionaries
- Java Mapping for Enumerations
- Java Mapping for Structures
- Java Mapping for Sequences
- Customizing the Java Mapping

Copyright 2016, ZeroC, Inc.

Java Mapping for Constants

Here are the sample constant definitions once more:

Slice

```
const bool      AppendByDefault = true;
const byte      LowerNibble = 0x0f;
const string    Advice = "Don't Panic!";
const short     TheAnswer = 42;
const double    PI = 3.1416;

enum Fruit { Apple, Pear, Orange };
const Fruit     FavoriteFruit = Pear;
```

Here are the generated definitions for these constants:

Java

```
public interface AppendByDefault {
    boolean value = true;
}

public interface LowerNibble {
    byte value = 15;
}

public interface Advice {
    String value = "Don't Panic!";
}

public interface TheAnswer {
    short value = 42;
}

public interface PI {
    double value = 3.1416;
}

public interface FavoriteFruit {
    Fruit value = Fruit.Pear;
}
```

As you can see, each Slice constant is mapped to a Java interface with the same name as the constant. The interface contains a member named `value` that holds the value of the constant.

Slice string literals that contain non-ASCII characters or universal character names are mapped to Java string literals with universal character names. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
const string Egg = "œuf";
const string Heart = "c\u0153ur";
const string Banana = "\U0001F34C";
```

is mapped to:

Java

```
public interface Egg {
    String value = "\u0153uf";
}

public interface Heart {
    String value = "c\u0153ur";
}

public interface Banana {
    String value = "\ud83c\udf4c";
}
```

See Also

- Constants and Literals
- Java Mapping for Identifiers
- Java Mapping for Modules
- Java Mapping for Built-In Types
- Java Mapping for Enumerations
- Java Mapping for Structures
- Java Mapping for Sequences
- Java Mapping for Dictionaries
- Java Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

Java Mapping for Exceptions

On this page:

- Java Mapping for User Exceptions
- Java Constructors for User Exceptions
- Java Mapping for Run-Time Exceptions

Java Mapping for User Exceptions

Here is a fragment of the Slice definition for our world time server once more:

Slice

```
exception GenericError {
    string reason;
};
exception BadTimeVal extends GenericError {};
exception BadZoneName extends GenericError {};
```

These exception definitions map as follows:

Java

```
public class GenericError extends Ice.UserException {
    public String reason;

    public GenericError() {}

    public GenericError(Throwable cause)
    {
        super(cause);
    }

    public GenericError(String reason)
    {
        this.reason = reason;
    }

    public GenericError(String reason, Throwable cause)
    {
        super(cause);
        this.reason = reason;
    }

    public String ice_name()
    {
        return "GenericError";
    }
}

public class BadTimeVal extends GenericError {
```


Copyright 2016, ZeroC, Inc.

```
    public BadTimeVal() {}

    public BadTimeVal(Throwable cause)
    {
        super(cause);
    }

    public BadTimeVal(String reason)
    {
        super(reason);
    }

    public BadTimeVal(String reason, Throwable cause)
    {
        super(reason, cause);
    }

    public String ice_name()
    {
        return "BadTimeVal";
    }
}

public class BadZoneName extends GenericError {
    public BadZoneName() {}

    public BadZoneName(Throwable cause)
    {
        super(cause);
    }

    public BadZoneName(String reason)
    {
        super(reason);
    }

    public BadZoneName(String reason, Throwable cause)
    {
        super(reason, cause);
    }

    public String ice_name()
    {
        return "BadZoneName";
```


Copyright 2016, ZeroC, Inc.

```
    }
}
```

Each Slice exception is mapped to a Java class with the same name. For each data member, the corresponding class contains a public data member. (Obviously, because `BadTimeVal` and `BadZoneName` do not have members, the generated classes for these exceptions also do not have members.) A JavaBean-style API is used for optional data members, and you can customize the mapping to force required members to use this same API.

The inheritance structure of the Slice exceptions is preserved for the generated classes, so `BadTimeVal` and `BadZoneName` inherit from `GenericError`.

Each exception also defines the `ice_name` member function, which returns the name of the exception.

All user exceptions are derived from the base class `Ice.UserException`. This allows you to catch all user exceptions generically by installing a handler for `Ice.UserException`. `Ice.UserException`, in turn, derives from `java.lang.Exception`.

`Ice.UserException` implements a `clone` method that is inherited by its derived exceptions, so you can make memberwise shallow copies of exceptions.

Note that the generated exception classes contain other member functions that are not shown. However, those member functions are internal to the Java mapping and are not meant to be called by application code.

Java Constructors for User Exceptions

Exceptions have a default constructor that initializes data members as follows:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	Null
`dictionary`	Null
`class/interface`	Null

The constructor won't explicitly initialize a data member if the default Java behavior for that type produces the desired results.

If you wish to ensure that data members of primitive and enumerated types are initialized to specific values, you can declare default values in your Slice definition. The default constructor initializes each of these data members to its declared value instead.

If an exception declares or inherits any data members, the generated class provides a second constructor that accepts one parameter for each data member so that you can construct and initialize an instance in a single statement (instead of first having to construct the instance and then assign to its members). For a derived exception, this constructor accepts one argument for each base exception member, plus one argument for each derived exception member, in base-to-derived order.

The generated class may include an additional constructor if the exception declares or inherits any optional data members.

The Slice compiler generates overloaded versions of all constructors that accept a trailing `Throwable` argument for preserving an exception chain.

Java Mapping for Run-Time Exceptions

The Ice run time throws run-time exceptions for a number of pre-defined error conditions. All run-time exceptions directly or indirectly derive from `Ice.LocalException` (which, in turn, derives from `java.lang.RuntimeException`).

`Ice.LocalException` implements a `clone` method that is inherited by its derived exceptions, so you can make memberwise shallow

Copyright 2016, ZeroC, Inc.

copies of exceptions.

Recall the inheritance diagram for user and run-time exceptions. By catching exceptions at the appropriate point in the hierarchy, you can handle exceptions according to the category of error they indicate:

- `Ice.LocalException`
  This is the root of the inheritance tree for run-time exceptions.
- `Ice.UserException`
  This is the root of the inheritance tree for user exceptions.
- `Ice.TimeoutException`
  This is the base exception for both operation-invocation and connection-establishment timeouts.
- `Ice.ConnectTimeoutException`
  This exception is raised when the initial attempt to establish a connection to a server times out.

For example, a `ConnectTimeoutException` can be handled as `ConnectTimeoutException`, `TimeoutException`, `LocalException`, or `java.lang.Exception`.

You will probably have little need to catch run-time exceptions as their most-derived type and instead catch them as `LocalException`; the fine-grained error handling offered by the remainder of the hierarchy is of interest mainly in the implementation of the Ice run time. Exceptions to this rule are the exceptions related to facet and object life cycles, which you may want to catch explicitly. These exceptions are `FacetNotExistException` and `ObjectNotExistException`, respectively.

See Also

- User Exceptions
- Run-Time Exceptions
- Java Mapping for Modules
- Java Mapping for Built-In Types
- Java Mapping for Enumerations
- Java Mapping for Structures
- Java Mapping for Sequences
- Java Mapping for Dictionaries
- Java Mapping for Constants
- Java Mapping for Optional Data Members
- JavaBean Mapping
- Versioning
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

Java Mapping for Interfaces

The mapping of Slice interfaces revolves around the idea that, to invoke a remote operation, you call a member function on a local class instance that is a proxy for the remote object. This makes the mapping easy and intuitive to use because making a remote procedure call is no different from making a local procedure call (apart from error semantics).

On this page:

- Java Classes Generated for an Interface
- Proxy Interfaces in Java
- The Ice.ObjectPrx Interface in Java
- Proxy Helpers in Java
- Using Proxy Methods in Java
- Object Identity and Proxy Comparison in Java
- Deserializing Proxies in Java

Java Classes Generated for an Interface

The compiler generates quite a few source files for each Slice interface. In general, for an interface `<interface-name>`, the following source files are created by the compiler:

- `<interface-name>.java`
  This source file declares the `<interface-name>` Java interface.
- `<interface-name>Holder.java`
  This source file defines a holder type for the interface.
- `<interface-name>Prx.java`
  This source file defines the proxy interface `<interface-name>Prx`.
- `<interface-name>PrxHelper.java`
  This source file defines the helper type for the interface's proxy.
- `<interface-name>PrxHolder.java`
  This source file defines the holder type for the interface's proxy.
- `_<interface-name>Operations.java`
  `_<interface-name>OperationsNC.java`
  These source files each define an interface that contains the operations corresponding to the Slice interface. The `Operations` interface is used for tie classes.

These are the files that contain code that is relevant to the client side. The compiler also generates a file that is specific to the server side, plus three additional files:

- `_<interface-name>Disp.java`
  This file contains the definition of the server-side skeleton class.
- `_<interface-name>Del.java`
- `_<interface-name>DelD.java`
- `_<interface-name>DelM.java`
  These files contain code that is internal to the Java mapping; they do not contain any functions of relevance to application programmers.

Proxy Interfaces in Java

On the client side, a Slice interface maps to a Java interface with methods that correspond to the operations on that interface. Consider the following simple interface:

Slice

```
interface Simple {
    void op();
};
```

The Slice compiler generates the following definition for use by the client:

Copyright 2016, ZeroC, Inc.

Java

```
public interface SimplePrx extends Ice.ObjectPrx {
    public void op();
    public void op(java.util.Map<String, String> __context);
}
```

As you can see, the compiler generates a *proxy interface* `SimplePrx`. In general, the generated name is <*interface-name*>`Prx`. If an interface is nested in a module `M`, the generated class is part of package `M`, so the fully-qualified name is `M.`<*interface-name*>`Prx`.

In the client's address space, an instance of `SimplePrx` is the local ambassador for a remote instance of the `Simple` interface in a server and is known as a proxy instance. All the details about the server-side object, such as its address, what protocol to use, and its object identity are encapsulated in that instance.

Note that `SimplePrx` inherits from `Ice.ObjectPrx`. This reflects the fact that all Ice interfaces implicitly inherit from `Ice::Object`.

For each operation in the interface, the proxy class has a method of the same name. For the preceding example, we find that the operation `op` has been mapped to the method `op`. Also note that `op` is overloaded: the second version of `op` has a parameter `__context` of type `java.util.Map<String, String>`. This parameter is for use by the Ice run time to store information about how to deliver a request. You normally do not need to use it. (We examine the `__context` parameter in detail in Request Contexts. The parameter is also used by IceStorm.)

Because all the <*interface-name*>`Prx` types are interfaces, you cannot instantiate an object of such a type. Instead, proxy instances are always instantiated on behalf of the client by the Ice run time, so client code never has any need to instantiate a proxy directly.The proxy references handed out by the Ice run time are always of type <*interface-name*>`Prx`; the concrete implementation of the interface is part of the Ice run time and does not concern application code.

A value of `null` denotes the null proxy. The null proxy is a dedicated value that indicates that a proxy points "nowhere" (denotes no object).

The `Ice.ObjectPrx` Interface in Java

All Ice objects have `Object` as the ultimate ancestor type, so all proxies inherit from `Ice.ObjectPrx`. `ObjectPrx` provides a number of methods:

Java

```
package Ice;

public interface ObjectPrx {
    boolean equals(java.lang.Object r);
    Identity ice_getIdentity();
    boolean ice_isA(String __id);
    boolean ice_isA(String __id, java.util.Map<String, String> ctx);
    String[] ice_ids();
    String[] ice_ids(java.util.Map<String, String> ctx);
    String ice_id();
    String ice_id(java.util.Map<String, String> ctx);
    void ice_ping();
    void ice_ping(java.util.Map<String, String> ctx);
    // ...
}
```

The methods behave as follows:

- **`equals`**

Copyright 2016, ZeroC, Inc.

This operation compares two proxies for equality. Note that all aspects of proxies are compared by this operation, such as the communication endpoints for the proxy. This means that, in general, if two proxies compare unequal, that does *not* imply that they denote different objects. For example, if two proxies denote the same Ice object via different transport endpoints, `equals` returns `false` even though the proxies denote the same object.

- **`ice_getIdentity`**
  This method returns the identity of the object denoted by the proxy. The identity of an Ice object has the following Slice type:

  Slice

```
module Ice {
    struct Identity {
        string name;
        string category;
    };
};
```

  To see whether two proxies denote the same object, first obtain the identity for each object and then compare the identities:

  Java

```
Ice.ObjectPrx o1 = ...;
Ice.ObjectPrx o2 = ...;
Ice.Identity i1 = o1.ice_getIdentity();
Ice.Identity i2 = o2.ice_getIdentity();

if (i1.equals(i2))
    // o1 and o2 denote the same object
else
    // o1 and o2 denote different objects
```

- **`ice_isA`**
  The `ice_isA` method determines whether the object denoted by the proxy supports a specific interface. The argument to `ice_isA` is a type ID. For example, to see whether a proxy of type `ObjectPrx` denotes a `Printer` object, we can write:

  Java

```
Ice.ObjectPrx o = ...;
if (o != null && o.ice_isA("::Printer"))
    // o denotes a Printer object
else
    // o denotes some other type of object
```

  Note that we are testing whether the proxy is null before attempting to invoke the `ice_isA` method. This avoids getting a `NullPointerException` if the proxy is null.

- **`ice_ids`**
  The `ice_ids` method returns an array of strings representing all of the type IDs that the object denoted by the proxy supports.

- **`ice_id`**
  The `ice_id` method returns the type ID of the object denoted by the proxy. Note that the type returned is the type of the actual object, which may be more derived than the static type of the proxy. For example, if we have a proxy of type `BasePrx`, with a static type ID of `::Base`, the return value of `ice_id` might be `::Base`, or it might something more derived, such as `::Derived`.

Copyright 2016, ZeroC, Inc.

- **ice_ping**
  The ice_ping method provides a basic reachability test for the object. If the object can physically be contacted (that is, the object exists and its server is running and reachable), the call completes normally; otherwise, it throws an exception that indicates why the object could not be reached, such as ObjectNotExistException or ConnectTimeoutException.

The ice_isA, ice_ids, ice_id, and ice_ping methods are remote operations and therefore support an additional overloading that accepts a request context. Also note that there are other methods in ObjectPrx, not shown here. These methods provide different ways to dispatch a call and also provide access to an object's facets.

Proxy Helpers in Java

For each Slice interface, apart from the proxy interface, the Slice-to-Java compiler creates a helper class: for an interface Simple, the name of the generated helper class is SimplePrxHelper. The helper class provides methods that support down-casting and type discovery:

Java

```
public final class SimplePrxHelper
extends Ice.ObjectPrxHelper implements SimplePrx {
    public static SimplePrx checkedCast(Ice.ObjectPrx b) {
        // ...
    }

    public static SimplePrx checkedCast(Ice.ObjectPrx b, Ice.Context ctx) {
        // ...
    }

    public static SimplePrx uncheckedCast(Ice.ObjectPrx b) {
        // ...
    }

    public static String ice_staticId() {
        // ...
    }

    // ...
}
```

For checkedCast, if the passed proxy is for an object of type Simple, or a proxy for an object with a type derived from Simple, the cast returns a non-null reference to a proxy of type SimplePrx; otherwise, if the passed proxy denotes an object of a different type (or if the passed proxy is null), the cast returns a null reference.

Given a proxy of any type, you can use a checkedCast to determine whether the corresponding object supports a given type, for example:

Java

```
Ice.ObjectPrx obj = ...;        // Get a proxy from somewhere...

SimplePrx simple = SimplePrxHelper.checkedCast(obj);
if (simple != null)
    // Object supports the Simple interface...
else
    // Object is not of type Simple...
```

Note that a checkedCast contacts the server. This is necessary because only the implementation of an object in the server has definite knowledge of the type of an object. As a result, a checkedCast may throw a ConnectTimeoutException or an ObjectNotExistExce

Copyright 2016, ZeroC, Inc.

ption. (This also explains the need for the helper class: the Ice run time must contact the server, so we cannot use a Java down-cast.)

In contrast, an uncheckedCast does not contact the server and unconditionally returns a proxy of the requested type. However, if you do use an uncheckedCast, you must be certain that the proxy really does support the type you are casting to; otherwise, if you get it wrong, you will most likely get a run-time exception when you invoke an operation on the proxy. The most likely error for such a type mismatch is OperationNotExistException. However, other exceptions, such as a marshaling exception are possible as well. And, if the object happens to have an operation with the correct name, but different parameter types, no exception may be reported at all and you simply end up sending the invocation to an object of the wrong type; that object may do rather nonsensical things. To illustrate this, consider the following two interfaces:

Slice

```
interface Process {
    void launch(int stackSize, int dataSize);
};

// ...

interface Rocket {
    void launch(float xCoord, float yCoord);
};
```

Suppose you expect to receive a proxy for a Process object and use an uncheckedCast to down-cast the proxy:

Java

```
Ice.ObjectPrx obj = ...;                    // Get proxy...
ProcessPrx process = ProcessPrxHelper.uncheckedCast(obj);  // No worries...
process.launch(40, 60);                     // Oops...
```

If the proxy you received actually denotes a Rocket object, the error will go undetected by the Ice run time: because int and float have the same size and because the Ice protocol does not tag data with its type on the wire, the implementation of Rocket::launch will simply misinterpret the passed integers as floating-point numbers.

In fairness, this example is somewhat contrived. For such a mistake to go unnoticed at run time, both objects must have an operation with the same name and, in addition, the run-time arguments passed to the operation must have a total marshaled size that matches the number of bytes that are expected by the unmarshaling code on the server side. In practice, this is extremely rare and an incorrect uncheckedCast typically results in a run-time exception.

A final warning about down-casts: you must use either a checkedCast or an uncheckedCast to down-cast a proxy. If you use a Java cast, the behavior is undefined.

Another method defined by every helper class is ice_staticId, which returns the type ID string corresponding to the interface. As an example, for the Slice interface Simple in module M, the string returned by ice_staticId is "::M::Simple".

Using Proxy Methods in Java

The base proxy class ObjectPrx supports a variety of methods for customizing a proxy. Since proxies are immutable, each of these "factory methods" returns a copy of the original proxy that contains the desired modification. For example, you can obtain a proxy configured with a ten second timeout as shown below:

Java

```
Ice.ObjectPrx proxy = communicator.stringToProxy(...);
proxy = proxy.ice_timeout(10000);
```


Copyright 2016, ZeroC, Inc.

A factory method returns a new proxy object if the requested modification differs from the current proxy, otherwise it returns the current proxy. With few exceptions, factory methods return a proxy of the same type as the current proxy, therefore it is generally not necessary to repeat a `checkedCast` or `uncheckedCast` after using a factory method. However, a regular cast is still required, as shown in the example below:

Java

```
Ice.ObjectPrx base = communicator.stringToProxy(...);
HelloPrx hello = HelloPrxHelper.checkedCast(base);
hello = (HelloPrx)hello.ice_timeout(10000); # Type is preserved
hello.sayHello();
```

The only exceptions are the factory methods `ice_facet` and `ice_identity`. Calls to either of these methods may produce a proxy for an object of an unrelated type, therefore they return a base proxy that you must subsequently down-cast to an appropriate type.

Object Identity and Proxy Comparison in Java

Proxies provide an `equals` method that compares proxies:

Java

```
interface ObjectPrx {
    boolean equals(java.lang.Object r);
}
```

Note that proxy comparison with `equals` uses *all* of the information in a proxy for the comparison. This means that not only the object identity must match for a comparison to succeed, but other details inside the proxy, such as the protocol and endpoint information, must be the same. In other words, comparison with `equals` tests for *proxy* identity, *not* object identity. A common mistake is to write code along the following lines:

Java

```
Ice.ObjectPrx p1 = ...;        // Get a proxy...
Ice.ObjectPrx p2 = ...;        // Get another proxy...

if (!p1.equals(p2)) {
    // p1 and p2 denote different objects       // WRONG!
} else {
    // p1 and p2 denote the same object         // Correct
}
```

Even though `p1` and `p2` differ, they may denote the same Ice object. This can happen because, for example, both `p1` and `p2` embed the same object identity, but each use a different protocol to contact the target object. Similarly, the protocols may be the same, but denote different endpoints (because a single Ice object can be contacted via several different transport endpoints). In other words, if two proxies compare equal with `equals`, we know that the two proxies denote the same object (because they are identical in all respects); however, if two proxies compare unequal with `equals`, we know absolutely nothing: the proxies may or may not denote the same object.

To compare the object identities of two proxies, you can use a helper function in the `Ice.Util` class:

Copyright 2016, ZeroC, Inc.

Java

```
package Ice;

public final class Util {
    public static int proxyIdentityCompare(ObjectPrx lhs, ObjectPrx rhs);
    public static int proxyIdentityAndFacetCompare(ObjectPrx lhs, ObjectPrx rhs);
    // ...
}
```

`proxyIdentityCompare` allows you to correctly compare proxies for identity:

Java

```
Ice.ObjectPrx p1 = ...;        // Get a proxy...
Ice.ObjectPrx p2 = ...;        // Get another proxy...

if (Ice.Util.proxyIdentityCompare(p1, p2) != 0) {
    // p1 and p2 denote different objects       // Correct
} else {
    // p1 and p2 denote the same object         // Correct
}
```

The function returns 0 if the identities are equal, `-1` if `p1` is less than `p2`, and 1 if `p1` is greater than `p2`. (The comparison uses `name` as the major and `category` as the minor sort key.)

The `proxyIdentityAndFacetCompare` function behaves similarly, but compares both the identity and the facet name.

In addition, the Java mapping provides two wrapper classes that allow you to wrap a proxy for use as the key of a hashed collection:

Java

```
package Ice;

public class ProxyIdentityKey {
    public ProxyIdentityKey(Ice.ObjectPrx proxy);
    public int hashCode();
    public boolean equals(java.lang.Object obj);
    public Ice.ObjectPrx getProxy();
}

public class ProxyIdentityFacetKey {
    public ProxyIdentityFacetKey(Ice.ObjectPrx proxy);
    public int hashCode();
    public boolean equals(java.lang.Object obj);
    public Ice.ObjectPrx getProxy();
}
```

The constructor caches the identity and the hash code of the passed proxy, so calls to `hashCode` and `equals` can be evaluated efficiently. The `getProxy` method returns the proxy that was passed to the constructor.

Copyright 2016, ZeroC, Inc.

As for the comparison functions, `ProxyIdentityKey` only uses the proxy's identity, whereas `ProxyIdentityFacetKey` also includes the facet name.

Deserializing Proxies in Java

Proxy objects implement the `java.io.Serializable` interface that enables serialization of proxies to and from a byte stream. You can use the standard class `java.io.ObjectInputStream` to deserialize all Slice types *except* proxies; proxies are a special case because they must be created by a communicator.

To supply a communicator for use in deserializing proxies, an application must use the class `Ice.ObjectInputStream`:

Java

```
package Ice;

public class ObjectInputStream extends java.io.ObjectInputStream
{
    public ObjectInputStream(Communicator communicator, java.io.InputStream stream)
        throws java.io.IOException;

    public Communicator getCommunicator();
}
```

The code shown below demonstrates how to use this class:

Java

```
Ice.Communicator communicator = ...
byte[] bytes = ... // data to be deserialized
java.io.ByteArrayInputStream byteStream = new java.io.ByteArrayInputStream(bytes);
Ice.ObjectInputStream in = new Ice.ObjectInputStream(communicator, byteStream);
Ice.ObjectPrx proxy = (Ice.ObjectPrx)in.readObject();
```

Ice raises `java.io.IOException` if an application attempts to deserialize a proxy without supplying a communicator.

See Also

- Interfaces, Operations, and Exceptions
- Proxies for Ice Objects
- Type IDs
- Java Mapping for Operations
- Request Contexts
- Operations on Object
- Proxy Methods
- Versioning
- IceStorm

Copyright 2016, ZeroC, Inc.

Java Mapping for Operations

On this page:

- Basic Java Mapping for Operations
- Normal and idempotent Operations in Java
- Passing Parameters in Java
  - In-Parameters in Java
  - Out-Parameters in Java
  - Null Parameters in Java
  - Optional Parameters in Java
- Exception Handling in Java
  - Exceptions and Out-Parameters

Basic Java Mapping for Operations

As we saw in the mapping for interfaces, for each operation on an interface, the proxy class contains a corresponding member function with the same name. To invoke an operation, you call it via the proxy. For example, here is part of the definitions for our file system:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };
    // ...
};
```

The `name` operation returns a value of type `string`. Given a proxy to an object of type `Node`, the client can invoke the operation as follows:

Java

```
NodePrx node = ...;             // Initialize proxy
String name = node.name();      // Get name via RPC
```

This illustrates the typical pattern for receiving return values: return values are returned by reference for complex types, and by value for simple types (such as `int` or `double`).

Normal and `idempotent` Operations in Java

You can add an `idempotent` qualifier to a Slice operation. As far as the signature for the corresponding proxy method is concerned, `idempotent` has no effect. For example, consider the following interface:

Slice

```
interface Example {
                string op1();
    idempotent  string op2();
};
```

The proxy interface for this is:

Copyright 2016, ZeroC, Inc.

Java

```
public interface ExamplePrx extends Ice.ObjectPrx {
    public String op1();
    public String op2();
}
```

Because `idempotent` affects an aspect of call dispatch, not interface, it makes sense for the two methods to be mapped the same.

Passing Parameters in Java

In-Parameters in Java

The parameter passing rules for the Java mapping are very simple: parameters are passed either by value (for simple types) or by reference (for complex types and type `string`). Semantically, the two ways of passing parameters are identical: it is guaranteed that the value of a parameter will not be changed by the invocation.

Here is an interface with operations that pass parameters of various types from client to server:

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ClientToServer {
    void op1(int i, float f, bool b, string s);
    void op2(NumberAndString ns, StringSeq ss, StringTable st);
    void op3(ClientToServer* proxy);
};
```

The Slice compiler generates the following proxy for these definitions:

Java

```
public interface ClientToServerPrx extends Ice.ObjectPrx {
    public void op1(int i, float f, boolean b, String s);
    public void op2(NumberAndString ns, String[] ss, java.util.Map st);
    public void op3(ClientToServerPrx proxy);
}
```

Given a proxy to a `ClientToServer` interface, the client code can pass parameters as in the following example:

Copyright 2016, ZeroC, Inc.

Java

```
ClientToServerPrx p = ...;             // Get proxy...

p.op1(42, 3.14f, true, "Hello world!"); // Pass simple literals

int i = 42;
float f = 3.14f;
boolean b = true;
String s = "Hello world!";
p.op1(i, f, b, s);                     // Pass simple variables

NumberAndString ns = new NumberAndString();
ns.x = 42;
ns.str = "The Answer";
String[] ss = { "Hello world!" };
java.util.HashMap st = new java.util.HashMap();
st.put(new Long(0), ns);
p.op2(ns, ss, st);                     // Pass complex variables

p.op3(p);                              // Pass proxy
```

Out-Parameters in Java

Java does not have pass-by-reference: parameters are always passed by value. For a function to modify one of its arguments, we must pass a reference (by value) to an object; the called function can then modify the object's contents via the passed reference.

To permit the called function to modify a parameter, the Java mapping uses *holder* classes. For example, for each of the built-in Slice types, such as `int` and `string`, the `Ice` package contains a corresponding holder class. Here are the definitions for the holder classes `Ice.IntHolder` and `Ice.StringHolder`:

Copyright 2016, ZeroC, Inc.

Java

```
package Ice;

public final class IntHolder {
    public IntHolder() {}
    public IntHolder(int value) {
        this.value = value;
    }
    public int value;
}

public final class StringHolder {
    public StringHolder() {}
    public StringHolder(String value) {
        this.value = value;
    }
    public String value;
}
```

A holder class has a public `value` member that stores the value of the parameter; the called function can modify the value by assigning to that member. The class also has a default constructor and a constructor that accepts an initial value.

For user-defined types, such as structures, the Slice-to-Java compiler generates a corresponding holder type. For example, here is the generated holder type for the `NumberAndString` structure we defined earlier:

Java

```
public final class NumberAndStringHolder {
    public NumberAndStringHolder() {}

    public NumberAndStringHolder(NumberAndString value) {
        this.value = value;
    }

    public NumberAndString value;
}
```

This looks exactly like the holder classes for the built-in types: we get a default constructor, a constructor that accepts an initial value, and the public `value` member.

Note that holder classes are generated for *every* Slice type you define. For example, for sequences, such as the `FruitPlatter` sequence, the compiler does not generate a special Java `FruitPlatter` type because sequences map to Java arrays. However, the compiler does generate a `FruitPlatterHolder` class, so we can pass a `FruitPlatter` array as an out-parameter.

To pass an out-parameter to an operation, we simply pass an instance of a holder class and examine the `value` member of each out-parameter when the call completes. Here are the same Slice definitions we saw earlier, but this time with all parameters being passed in the `out` direction:

Copyright 2016, ZeroC, Inc.

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ServerToClient {
    void op1(out int i, out float f, out bool b, out string s);
    void op2(out NumberAndString ns,
             out StringSeq ss,
             out StringTable st);
    void op3(out ServerToClient* proxy);
};
```

The Slice compiler generates the following code for these definitions:

Java

```
public interface ClientToServerPrx extends Ice.ObjectPrx {
    public void op1(Ice.IntHolder i, Ice.FloatHolder f,
                    Ice.BooleanHolder b, Ice.StringHolder s);
    public void op2(NumberAndStringHolder ns,
                    StringSeqHolder ss, StringTableHolder st);
    public void op3(ClientToServerPrxHolder proxy);
}
```

Given a proxy to a `ServerToClient` interface, the client code can pass parameters as in the following example:

Copyright 2016, ZeroC, Inc.

Java

```
ClientToServerPrx p = ...;             // Get proxy...

Ice.IntHolder ih = new Ice.IntHolder();
Ice.FloatHolder fh = new Ice.FloatHolder();
Ice.BooleanHolder bh = new Ice.BooleanHolder();
Ice.StringHolder sh = new Ice.StringHolder();
p.op1(ih, fh, bh, sh);

NumberAndStringHolder nsh = new NumberAndString();
StringSeqHolder ssh = new StringSeqHolder();
StringTableHolder sth = new StringTableHolder();
p.op2(nsh, ssh, sth);

ServerToClientPrxHolder stcph = new ServerToClientPrxHolder();
p.op3(stch);

System.out.writeln(ih.value);   // Show one of the values
```

Again, there are no surprises in this code: the various holder instances contain values once the operation invocation completes and the `value` member of each instance provides access to those values.

Null Parameters in Java

Some Slice types naturally have "empty" or "not there" semantics. Specifically, sequences, dictionaries, and strings all can be `null`, but the corresponding Slice types do not have the concept of a null value. To make life with these types easier, whenever you pass `null` as a parameter or return value of type sequence, dictionary, or string, the Ice run time automatically sends an empty sequence, dictionary, or string to the receiver.

This behavior is useful as a convenience feature: especially for deeply-nested data types, members that are sequences, dictionaries, or strings automatically arrive as an empty value at the receiving end. This saves you having to explicitly initialize, for example, every string element in a large sequence before sending the sequence in order to avoid `NullPointerException`. Note that using null parameters in this way does *not* create null semantics for Slice sequences, dictionaries, or strings. As far as the object model is concerned, these do not exist (only *empty* sequences, dictionaries, and strings do). For example, whether you send a string as `null` or as an empty string makes no difference to the receiver: either way, the receiver sees an empty string.

Optional Parameters in Java

The Java mapping uses the generic class `Ice.Optional` to encapsulate optional parameters. All Slice types (including `string`) that map to Java reference types use the `Ice.Optional` generic class. To minimize garbage, the mapping also defines several non-generic classes to hold the unboxed primitive types `boolean`, `byte`, `short`, `int`, `long`, `float`, and `double`. For example, an optional `boolean` parameter maps to the `Ice.BooleanOptional` class.

Optional return values and output parameters are mapped to instances of `Ice.Optional` or one of the primitive wrapper classes, depending on their types. For operations with optional *input* parameters, the proxy provides a set of overloaded methods that accept them as optional values, and another set of methods that accept them as required values. Consider the following operation:

Copyright 2016, ZeroC, Inc.

Slice

```
optional(1) int execute(optional(2) string params, out optional(3) float value);
```

The default Java mapping for this operation is shown below:

Java

```
Ice.IntOptional execute(String params, Ice.FloatOptional value, ...);
Ice.IntOptional execute(Ice.Optional<String> params, Ice.FloatOptional value, ...);
```

For cases where you are passing values for all optional input parameters, it is more efficient to use the required mapping and avoid creating temporary optional values.

A client can invoke `execute` as shown below:

Java

```
Ice.IntOptional i;
Ice.FloatOptional v = new Ice.FloatOptional();

i = proxy.execute("--file log.txt", v);                          // required
mapping
i = proxy.execute(new Ice.Optional<String>("--file log.txt"), v); // optional
mapping
i = proxy.execute((Ice.Optional<String>)null, v);                // params is
unset

if(v.isSet())
    System.out.println("value = " + v.get());
```

Passing `null` where an `Ice.Optional` value is expected is equivalent to passing an `Ice.Optional` instance whose value is unset. For optional parameters of reference types, passing an explicit `null` requires a cast as shown in the last call to `execute` above.

Whereas the mapping for required output parameters uses `Holder` classes to transfer their values from the Ice run time to the caller, the optional mapping uses `Optional` values. When invoking a proxy operation, the client must pass an instance of the appropriate `Optional` type for each output parameter. The Ice run time invokes `clear` on the client's `Optional` value if the server does not supply a value for an output parameter, therefore it is safe to pass an `Optional` instance that already has a value.

A well-behaved program must not assume that an optional parameter always has a value. Calling `get` on an `Optional` instance for which no value is set raises `java.lang.IllegalStateException`.

Exception Handling in Java

Any operation invocation may throw a run-time exception and, if the operation has an exception specification, may also throw user exceptions. Suppose we have the following simple interface:

Copyright 2016, ZeroC, Inc.

Slice

```
exception Tantrum {
    string reason;
};

interface Child {
    void askToCleanUp() throws Tantrum;
};
```

Slice exceptions are thrown as Java exceptions, so you can simply enclose one or more operation invocations in a `try-catch` block:

Java

```
ChildPrx child = ...;   // Get child proxy...

try {
    child.askToCleanUp();
} catch (Tantrum t) {
    System.out.write("The child says: ");
    System.out.writeln(t.reason);
}
```

Typically, you will catch only a few exceptions of specific interest around an operation invocation; other exceptions, such as unexpected run-time errors, will typically be handled by exception handlers higher in the hierarchy. For example:

Copyright 2016, ZeroC, Inc.

Java

```
public class Client {
    static void run() {
        ChildPrx child = ...;   // Get child proxy...
        try {
            child.askToCleanUp();
        } catch (Tantrum t) {
            System.out.print("The child says: ");
            System.out.println(t.reason);
            child.scold();          // Recover from error...
        }
        child.praise();             // Give positive feedback...
    }

    public static void
    main(String[] args)
    {
        try {
            // ...
            run();
            // ...
        } catch (Ice.LocalException e) {
            e.printStackTrace();
        } catch (Ice.UserException e) {
            System.err.println(e.getMessage());
        }
    }
}
```

This code handles a specific exception of local interest at the point of call and deals with other exceptions generically. (This is also the strategy we used for our first simple application.)

Exceptions and Out-Parameters

The Ice run time makes no guarantees about the state of out-parameters when an operation throws an exception: the parameter may still have its original value or may have been changed by the operation's implementation in the target object. In other words, for out-parameters, Ice provides the weak exception guarantee [1] but does not provide the strong exception guarantee.

This is done for reasons of efficiency: providing the strong exception guarantee would require more overhead than can be justified.

See Also

- Operations
- Java Mapping for Exceptions
- Java Mapping for Sequences
- Java Mapping for Interfaces
- Java Mapping for Optional Data Members
- Collocated Invocation and Dispatch

References

1. Sutter, H. 1999. Exceptional C++: 47 Engineering Puzzles, Programming Problems, and Solutions. Reading, MA: Addison-Wesley.

Copyright 2016, ZeroC, Inc.

Java Mapping for Classes

On this page:

- Basic Java Mapping for Classes
- Operations Interfaces in Java
- Inheritance from Ice.Object in Java
- Class Data Members in Java
- Class Operations in Java
- Class Factories in Java
- Class Constructors in Java

Basic Java Mapping for Classes

A Slice class is mapped to a Java class with the same name. The generated class contains a public data member for each Slice data member (just as for structures and exceptions), and a member function for each operation. Consider the following class definition:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Slice compiler generates the following code for this definition:

Java

```
public interface _TimeOfDayOperations {
    String format(Ice.Current current);
}

public interface _TimeOfDayOperationsNC {
    String format();
}

public abstract class TimeOfDay extends Ice.ObjectImpl
                                implements _TimeOfDayOperations,
_TimeOfDayOperationsNC {
    public short hour;
    public short minute;
    public short second;

    public TimeOfDay();
    public TimeOfDay(short hour, short minute, short second);
    // ...
}
```

There are a number of things to note about the generated code:

1. The compiler generates "operations interfaces" called `_TimeOfDayOperations` and `_TimeOfDayOperationsNC`. These

Copyright 2016, ZeroC, Inc.

interfaces contain a method for each Slice operation of the class.
2. The generated class `TimeOfDay` inherits (indirectly) from `Ice.Object`. This means that all classes implicitly inherit from `Ice.Object`, which is the ultimate ancestor of all classes. Note that `Ice.Object` is *not* the same as `Ice.ObjectPrx`. In other words, you *cannot* pass a class where a proxy is expected and vice versa.
   If a class has only data members, but no operations, the compiler generates a non-abstract class.
3. The generated class contains a public member for each Slice data member.
4. The generated class inherits member functions for each Slice operation from the operations interfaces.
5. The generated class contains two constructors.

There is quite a bit to discuss here, so we will look at each item in turn.

Operations Interfaces in Java

The methods in the `_<interface-name>Operations` interface have an additional trailing parameter of type `Ice.Current`, whereas the methods in the `_<interface-name>OperationsNC` interface lack this additional trailing parameter. The methods without the `Current` parameter simply forward to the methods with a `Current` parameter, supplying a default `Current`. For now, you can ignore this parameter and pretend it does not exist.

If a class has only data members, but no operations, the compiler omits generating the `<_interface-name>Operations` and `_<interface-name>OperationsNC` interfaces.

Inheritance from `Ice.Object` in Java

Like interfaces, classes implicitly inherit from a common base class, `Ice.Object`. However, as shown in the illustration below, classes inherit from `Ice.Object` instead of `Ice.ObjectPrx` (which is at the base of the inheritance hierarchy for proxies). As a result, you cannot pass a class where a proxy is expected (and vice versa) because the base types for classes and proxies are not compatible.

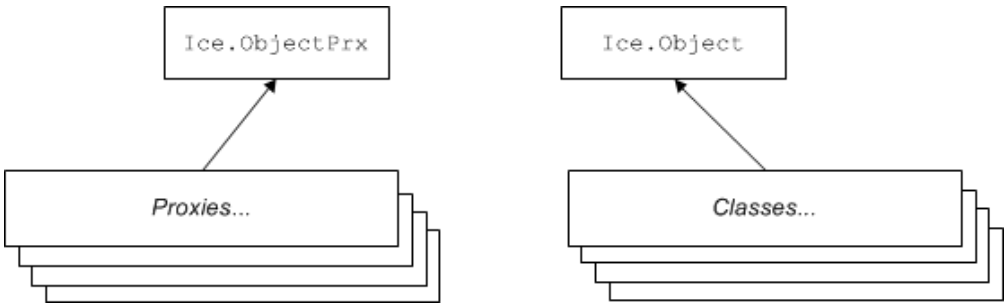

Inheritance from `Ice.ObjectPrx` and `Ice.Object`.

`Ice.Object` contains a number of member functions:

Copyright 2016, ZeroC, Inc.

Java

```
package Ice;

public interface Object
{
    boolean ice_isA(String s);
    boolean ice_isA(String s, Current current);

    void ice_ping();
    void ice_ping(Current current);

    String[] ice_ids();
    String[] ice_ids(Current current);

    String ice_id();
    String ice_id(Current current);

    void ice_preMarshal();
    void ice_postUnmarshal();

    DispatchStatus ice_dispatch(Request request,
DispatchInterceptorAsyncCallback cb);
}
```

The member functions of `Ice.Object` behave as follows:

- `ice_isA`
  This function returns `true` if the object supports the given type ID, and `false` otherwise.
- `ice_ping`
  As for interfaces, `ice_ping` provides a basic reachability test for the class.
- `ice_ids`
  This function returns a string sequence representing all of the type IDs supported by this object, including `::Ice::Object`.
- `ice_id`
  This function returns the actual run-time type ID for a class. If you call `ice_id` through a reference to a base instance, the returned type id is the actual (possibly more derived) type ID of the instance.
- `ice_preMarshal`
  The Ice run time invokes this function prior to marshaling the object's state, providing the opportunity for a subclass to validate its declared data members.
- `ice_postUnmarshal`
  The Ice run time invokes this function after unmarshaling an object's state. A subclass typically overrides this function when it needs to perform additional initialization using the values of its declared data members.
- `ice_dispatch`
  This function dispatches an incoming request to a servant. It is used in the implementation of dispatch interceptors.

Note that the generated class does not override `hashCode` and `equals`. This means that classes are compared using shallow reference equality, not value equality (as is used for structures).

All Slice classes derive from `Ice.Object` via the `Ice.ObjectImpl` abstract base class. `ObjectImpl` implements the `java.io.Serializable` interface to support Java's serialization facility. `ObjectImpl` also supplies an implementation of `clone` that returns a shallow memberwise copy.

Class Data Members in Java

Copyright 2016, ZeroC, Inc.

By default, data members of classes are mapped exactly as for structures and exceptions: for each data member in the Slice definition, the generated class contains a corresponding public data member. A JavaBean-style API is used for optional data members, and you can customize the mapping to force required members to use this same API.

If you wish to restrict access to a data member, you can modify its visibility using the `protected` metadata directive. The presence of this directive causes the Slice compiler to generate the data member with protected visibility. As a result, the member can be accessed only by the class itself or by one of its subclasses. For example, the `TimeOfDay` class shown below has the `protected` metadata directive applied to each of its data members:

Slice

```
class TimeOfDay {
    ["protected"] short hour;   // 0 - 23
    ["protected"] short minute; // 0 - 59
    ["protected"] short second; // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Slice compiler produces the following generated code for this definition:

Java

```
public abstract class TimeOfDay extends Ice.ObjectImpl
                                implements _TimeOfDayOperations,
                                           _TimeOfDayOperationsNC {
    protected short hour;
    protected short minute;
    protected short second;

    public TimeOfDay();
    public TimeOfDay(short hour, short minute, short second);
    // ...
}
```

For a class in which all of the data members are protected, the metadata directive can be applied to the class itself rather than to each member individually. For example, we can rewrite the `TimeOfDay` class as follows:

Slice

```
["protected"] class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

Note that you can optionally customize the mapping for data members to use getters and setters instead.

Copyright 2016, ZeroC, Inc.

Class Operations in Java

Operations of classes are mapped to abstract member functions in the generated class. This means that, if a class contains operations (such as the `format` operation of our `TimeOfDay` class), you must provide an implementation of the operation in a class that is derived from the generated class. For example:

Java

```
public class TimeOfDayI extends TimeOfDay {
    public String format(Ice.Current current) {
        DecimalFormat df = (DecimalFormat)DecimalFormat.getInstance();
        df.setMinimumIntegerDigits(2);
        return new String(df.format(hour) + ":" + df.format(minute) + ":" +
                          df.format(second));
    }
}
```

Class Factories in Java

Having created a class such as `TimeOfDayI`, we have an implementation and we can instantiate the `TimeOfDayI` class, but we cannot receive it as the return value or as an out-parameter from an operation invocation. To see why, consider the following simple interface:

Slice

```
interface Time {
    TimeOfDay get();
};
```

When a client invokes the `get` operation, the Ice run time must instantiate and return an instance of the `TimeOfDay` class. However, `TimeOfDay` is an abstract class that cannot be instantiated. Unless we tell it, the Ice run time cannot magically know that we have created a `TimeOfDayI` class that implements the abstract `format` operation of the `TimeOfDay` abstract class. In other words, we must provide the Ice run time with a factory that knows that the `TimeOfDay` abstract class has a `TimeOfDayI` concrete implementation. The `Ice::Communicator` interface provides us with the necessary operations:

Slice

```
module Ice {
    local interface ObjectFactory {
        Object create(string type);
        void destroy();
    };

    local interface Communicator {
        void addObjectFactory(ObjectFactory factory, string id);
        ObjectFactory findObjectFactory(string id);
        // ...
    };
};
```


Copyright 2016, ZeroC, Inc.

To supply the Ice run time with a factory for our `TimeOfDayI` class, we must implement the `ObjectFactory` interface:

Java

```
class ObjectFactory implements Ice.ObjectFactory {
    public Ice.Object create(String type) {
        if (type.equals(M.TimeOfDay.ice_staticId())) {
            return new TimeOfDayI();
        }
        assert(false);
        return null;
    }

    public void destroy() {
        // Nothing to do
    }
}
```

The object factory's `create` method is called by the Ice run time when it needs to instantiate a `TimeOfDay` class. The factory's `destroy` m ethod is called by the Ice run time when its communicator is destroyed.

The `create` method is passed the type ID of the class to instantiate. For our `TimeOfDay` class, the type ID is `"::M::TimeOfDay"`. Our implementation of `create` checks the type ID: if it matches, the method instantiates and returns a `TimeOfDayI` object. For other type IDs, the method asserts because it does not know how to instantiate other types of objects.

Note that we used the `ice_staticId` method to obtain the type ID rather than embedding a literal string. Using a literal type ID string in your code is discouraged because it can lead to errors that are only detected at run time. For example, if a Slice class or one of its enclosing modules is renamed and the literal string is not changed accordingly, a receiver will fail to unmarshal the object and the Ice run time will raise `NoObjectFactoryException`. By using `ice_staticId` instead, we avoid any risk of a misspelled or obsolete type ID, and we can discover at compile time if a Slice class or module has been renamed.

Given a factory implementation, such as our `ObjectFactory`, we must inform the Ice run time of the existence of the factory:

Java

```
Ice.Communicator ic = ...;
ic.addObjectFactory(new ObjectFactory(), M.TimeOfDay.ice_staticId());
```

Now, whenever the Ice run time needs to instantiate a class with the type ID `"::M::TimeOfDay"`, it calls the `create` method of the registered `ObjectFactory` instance.

The `destroy` operation of the object factory is invoked by the Ice run time when the communicator is destroyed. This gives you a chance to clean up any resources that may be used by your factory. Do not call `destroy` on the factory while it is registered with the communicator — if you do, the Ice run time has no idea that this has happened and, depending on what your `destroy` implementation is doing, may cause undefined behavior when the Ice run time tries to next use the factory.

The run time guarantees that `destroy` will be the last call made on the factory, that is, `create` will not be called concurrently with `destroy` , and `create` will not be called once `destroy` has been called. However, calls to `create` can be made concurrently.

Note that you cannot register a factory for the same type ID twice: if you call `addObjectFactory` with a type ID for which a factory is registered, the Ice run time throws an `AlreadyRegisteredException`.

Finally, keep in mind that if a class has only data members, but no operations, you need not create and register an object factory to transmit instances of such a class. Only if a class has operations do you have to define and register an object factory.

Class Constructors in Java

Copyright 2016, ZeroC, Inc.

Classes have a default constructor that initializes data members as follows:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	Null
`dictionary`	Null
`class/interface`	Null

The constructor won't explicitly initialize a data member if the default Java behavior for that type produces the desired results.

If you wish to ensure that data members of primitive and enumerated types are initialized to specific values, you can declare default values in your Slice definition. The default constructor initializes each of these data members to its declared value instead.

The generated class also contains a second constructor that accepts one argument for each member of the class. This allows you to create and initialize a class in a single statement, for example:

Java

```
TimeOfDayI tod = new TimeOfDayI(14, 45, 00); // 14:45pm
```

For derived classes, the constructor requires an argument for every member of the class, including inherited members. For example, consider the the definition from Class Inheritance once more:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
};

class DateTime extends TimeOfDay {
    short day;          // 1 - 31
    short month;        // 1 - 12
    short year;         // 1753 onwards
};
```

The constructors for the generated classes are as follows:

Copyright 2016, ZeroC, Inc.

Java

```
public class TimeOfDay extends Ice.ObjectImpl {
    public TimeOfDay() {}

    public TimeOfDay(short hour, short minute, short second)
    {
        this.hour = hour;
        this.minute = minute;
        this.second = second;
    }

    // ...
}

public class DateTime extends TimeOfDay
{
    public DateTime()
    {
        super();
    }

    public DateTime(short hour, short minute, short second,
                    short day, short month, short year)
    {
        super(hour, minute, second);
        this.day = day;
        this.month = month;
        this.year = year;
    }

    // ...
}
```

If you want to instantiate and initialize a `DateTime` instance, you must either use the default constructor or provide values for all of the data members of the instance, including data members of any base classes.

See Also

- Classes
- Class Inheritance
- Java Mapping for Optional Data Members
- Type IDs
- Serializable Objects in Java
- JavaBean Mapping
- The Current Object
- Dispatch Interceptors

Copyright 2016, ZeroC, Inc.

Java Mapping for Optional Data Members

On this page:

- Overview of Java Mapping for Optional Data Members
- Java Helper Classes for Optional Values

Overview of Java Mapping for Optional Data Members

The Java mapping for optional data members in Slice classes and exceptions uses a JavaBean-style API that provides methods to get, set, and clear a member's value, and test whether a value is set. Consider the following Slice definition:

Slice

```
class C
{
    string name;
    optional(2) string alternateName;
    optional(5) bool active;
};
```

The generated Java code provides the following API:

Java

```
public class C ...
{
    public C();
    public C(String name);
    public C(String name, String alternateName, boolean active);

    public String name;

    public String getAlternateName();
    public void setAlternateName(String v);
    public boolean hasAlternateName();
    public void clearAlternateName();
    public void optionalAlternateName(Ice.Optional<String> v);
    public Ice.Optional<String> optionalAlternateName();

    public boolean getActive();
    public void setActive(boolean v);
    public boolean isActive();
    public boolean hasActive();
    public void clearActive();
    public void optionalActive(Ice.BooleanOptional v);
    public Ice.BooleanOptional optionalActive();

    ...
}
```

If a class or exception declares any required data members, the generated class includes an overloaded constructor that accepts values for

Copyright 2016, ZeroC, Inc.

just the required members; optional members remain unset unless their Slice definitions specify a default value. Another overloaded constructor accepts values for all data members.

The `has` method allows you to test whether a member's value has been set, and the `clear` method removes any existing value for a member.

Calling a `get` method when the member's value has not been set raises `java.lang.IllegalStateException`.

The `optional` methods provide an alternate API that uses an `Optional` object to encapsulate the value, as discussed below.

Java Helper Classes for Optional Values

Ice defines the following classes to encapsulate optional values of primitive types:

- `Ice.BooleanOptional`
- `Ice.ByteOptional`
- `Ice.DoubleOptional`
- `Ice.FloatOptional`
- `Ice.IntOptional`
- `Ice.LongOptional`
- `Ice.ShortOptional`

These classes all share the same API for getting, setting, and testing an optional value. We'll use the `IntOptional` class as an example:

Java

```
public class IntOptional
{
    public IntOptional();               // Value is unset
    public IntOptional(int v);          // Value is set
    public IntOptional(IntOptional opt); // Copies the state of the argument
    public int get();                   // Raises IllegalStateException if unset
    public void set(int v);             // Value is set
    public void set(IntOptional opt);   // Copies the state of the argument
    public boolean isSet();
    public void clear();

    ...
}
```

The `Ice.Optional` generic class encapsulates values of reference types and offers an identical API:

Copyright 2016, ZeroC, Inc.

Java

```
public class Optional<T>
{
    public Optional();                // Value is unset
    public Optional(T v);             // Value is set
    public Optional(Optional<T> opt); // Copies the state of the argument
    public T get();                   // Raises IllegalStateException if unset
    public void set(T v);             // Value is set
    public void set(Optional<T> opt); // Copies the state of the argument
    public boolean isSet();
    public void clear();


    public static <T> Optional<T> O(T v);
    public static BooleanOptional O(boolean v);
    public static ByteOptional O(byte v);
    public static ShortOptional O(short v);
    public static IntOptional O(int v);
    public static LongOptional O(long v);
    public static FloatOptional O(float v);
    public static DoubleOptional O(double v);

    ...
}
```

To improve the readability of code that creates optional values, the `Optional` class defines overloaded `O` methods that simplify the creation of these objects:

Java

```
import static Ice.Optional.O;

C obj = ...;
Ice.Optional<C> opt = O(obj);

Ice.IntOptional i = O(5);
```

See Also

- Optional Data Members

Copyright 2016, ZeroC, Inc.

Serializable Objects in Java

In Java terminology, a *serializable object* typically refers to an object that implements the `java.io.Serializable` interface and therefore supports serialization to and from a byte stream. All Java classes generated from Slice definitions implement the `java.io.Serializable` interface.

In addition to serializing Slice types, applications may also need to incorporate foreign types into their Slice definitions. Ice allows you to pass Java serializable objects directly as operation parameters or as fields of another data type. For example:

Slice

```
["java:serializable:SomePackage.JavaClass"]
sequence<byte> JavaObj;
struct MyStruct {
    int i;
    JavaObj o;
};

interface Example {
    void op(JavaObj inObj, MyStruct s, out JavaObj outObj);
};
```

The generated code for `MyStruct` contains a member `i` of type `int` and a member `o` of type `SomePackage.JavaClass`:

Java

```
public final class MyStruct implements java.lang.Cloneable {
    public int i;
    public SomePackage.JavaClass o;

    // ...
}
```

Similarly, the signature for `op` has parameters of type `JavaClass` and `MyStruct` for the in-parameters, and `Ice.Holder<SomePackage.JavaClass>` for the out-parameter. (Out-parameters are always passed as `Ice.Holder<`*`class`*`>`.)

Java

```
void op(SomePackage.JavaClass inObj,
        MyStruct s,
        Ice.Holder<SomePackage.JavaClass> outObj);
```

Of course, your client and server code must have an implementation of `JavaClass` that derives from `java.io.Serializable`:

Copyright 2016, ZeroC, Inc.

Java

```
package SomePackage;

public class JavaClass implements java.io.Serializable {
    // ...
}
```

You can implement this class in any way you see fit — the Ice run time does not place any other requirements on the implementation.

See Also

- Serializable Objects

Copyright 2016, ZeroC, Inc.

Customizing the Java Mapping

You can customize the code that the Slice-to-Java compiler produces by annotating your Slice definitions with metadata. This section describes how metadata influences several aspects of the generated Java code.

On this page:

- Java Packages
  - Java Package Configuration Properties
- Custom Types in Java
  - Metadata in Java
  - Defining a Custom Sequence Type in Java
  - Defining a Custom Dictionary Type in Java
  - Using Custom Type Metadata in Java
  - Mapping for Modified Out Parameters in Java
- Buffer Types in Java
- JavaBean Mapping
  - JavaBean Generated Methods
  - JavaBean Metadata
- Overriding serialVersionUID

Java Packages

By default, the scope of a Slice definition determines the package of its mapped Java construct. A Slice type defined in a module hierarchy is mapped to a type residing in the equivalent Java package.

There are times when applications require greater control over the packaging of generated Java classes. For instance, a company may have software development guidelines that require all Java classes to reside in a designated package. One way to satisfy this requirement is to modify the Slice module hierarchy so that the generated code uses the required package by default. In the example below, we have enclosed the original definition of `Workflow::Document` in the modules `com::acme` so that the compiler will create the class in the `com.acme` package:

Slice

```
module com {
    module acme {
        module Workflow {
            class Document {
                // ...
            };
        };
    };
};
```

There are two problems with this workaround:

1. It incorporates the requirements of an implementation language into the application's interface specification.
2. Developers using other languages, such as C++, are also affected.

The Slice-to-Java compiler provides a better way to control the packages of generated code through the use of global metadata. The example above can be converted as follows:

Copyright 2016, ZeroC, Inc.

Slice

```
[["java:package:com.acme"]]
module Workflow {
    class Document {
        // ...
    };
};
```

The global metadata directive `java:package:com.acme` instructs the compiler to generate all of the classes resulting from definitions in this Slice file into the Java package `com.acme`. The net effect is the same: the class for `Document` is generated in the package `com.acme.Workflow`. However, we have addressed the two shortcomings of the first solution by reducing our impact on the interface specification: the Slice-to-Java compiler recognizes the package metadata directive and modifies its actions accordingly, whereas the compilers for other language mappings simply ignore it.

Java Package Configuration Properties

Using global metadata to alter the default package of generated classes has ramifications for the Ice run time when unmarshaling exceptions and concrete class types. The Ice run time dynamically loads generated classes by translating their Slice type ids into Java class names. For example, the Ice run time translates the Slice type id `::Workflow::Document` into the class name `Workflow.Document`.

However, when the generated classes are placed in a user-specified package, the Ice run time can no longer rely on the direct translation of a Slice type id into a Java class name, and therefore requires additional configuration so that it can successfully locate the generated classes. Two configuration properties are supported:

- `Ice.Package.`*`Module=package`*
  Associates a top-level Slice module with the package in which it was generated.

  > Only top-level module names are allowed; the semantics of global metadata prevent a nested module from being generated into a different package than its enclosing module.

- `Ice.Default.Package=package`
  Specifies a default package to use if other attempts to load a class have failed.

The behavior of the Ice run time when unmarshaling an exception or concrete class is described below:

1. Translate the Slice type id into a Java class name and attempt to load the class.
2. If that fails, extract the top-level module from the type id and check for an `Ice.Package` property with a matching module name. If found, prepend the specified package to the class name and try to load the class again.
3. If that fails, check for the presence of `Ice.Default.Package`. If found, prepend the specified package to the class name and try to load the class again.
4. If the class still cannot be loaded, the instance may be sliced.

Continuing our example from the previous section, we can define the following property:

```
Ice.Package.Workflow=com.acme
```

Alternatively, we could achieve the same result with this property:

```
Ice.Default.Package=com.acme
```


Copyright 2016, ZeroC, Inc.

Custom Types in Java

One of the more powerful applications of metadata is the ability to tailor the Java mapping for sequence and dictionary types to match the needs of your application.

Metadata in Java

The metadata for specifying a custom type has the following format:

```
java:type:instance-type[:formal-type]
```

The formal type is optional; the compiler uses a default value if one is not defined. The instance type must satisfy an is-A relationship with the formal type: either the same class is specified for both types, or the instance type must be derived from the formal type.

The Slice-to-Java compiler generates code that uses the formal type for all occurrences of the modified Slice definition except when the generated code must instantiate the type, in which case the compiler uses the instance type instead.

The compiler performs no validation on your custom types. Misspellings and other errors will not be apparent until you compile the generated code.

Defining a Custom Sequence Type in Java

Although the default mapping of a sequence type to a native Java array is efficient and typesafe, it is not always the most convenient representation of your data. To use a different representation, specify the type information in a metadata directive, as shown in the following example:

Slice

```
["java:type:java.util.LinkedList<String>"]
sequence<string> StringList;
```

It is your responsibility to use a type parameter for the Java class (`String` in the example above) that is the correct mapping for the sequence's element type.

The compiler requires the formal type to implement `java.util.List<`*E*`>`, where *E* is the Java mapping of the element type. If you do not specify a formal type, the compiler uses `java.util.List<`*E*`>` by default.

Note that extra care must be taken when defining custom types that contain nested generic types, such as a custom sequence whose element type is also a custom sequence. The Java compiler strictly enforces type safety, therefore any compatibility issues in the custom type metadata will be apparent when the generated code is compiled.

Defining a Custom Dictionary Type in Java

The default instance type for a dictionary is `java.util.HashMap<`*K*`, `*V*`>`, where *K* is the Java mapping of the key type and *V* is the Java mapping of the value type. If the semantics of a `HashMap` are not suitable for your application, you can specify an alternate type using metadata as shown in the example below:

Copyright 2016, ZeroC, Inc.

Slice

```
["java:type:java.util.TreeMap<String, String>"]
dictionary<string, string> StringMap;
```

It is your responsibility to use type parameters for the Java class (`String` in the example above) that are the correct mappings for the dictionary's key and value types.

The compiler requires the formal type to implement `java.util.Map<`*K, V*`>`. If you do not specify a formal type, the compiler uses this type by default.

Note that extra care must be taken when defining dictionary types that contain nested generic types, such as a dictionary whose element type is a custom sequence. The Java compiler strictly enforces type safety, therefore any compatibility issues in the custom type metadata will be apparent when the generated code is compiled.

Using Custom Type Metadata in Java

You can define custom type metadata in a variety of situations. The simplest scenario is specifying the metadata at the point of definition:

Slice

```
["java:type:java.util.LinkedList<String>"]
sequence<string> StringList;
```

Defined in this manner, the Slice-to-Java compiler uses `java.util.List<String>` (the default formal type) for all occurrences of `String List`, and `java.util.LinkedList<String>` when it needs to instantiate `StringList`.

You may also specify a custom type more selectively by defining metadata for a data member, parameter or return value. For instance, the mapping for the original Slice definition might be sufficient in most situations, but a different mapping is more convenient in particular cases. The example below demonstrates how to override the sequence mapping for the data member of a structure as well as for several operations:

Slice

```
sequence<string> StringSeq;

struct S {
    ["java:type:java.util.LinkedList<String>"] StringSeq seq;
};

interface I {
    ["java:type:java.util.ArrayList<String>"] StringSeq
    modifiedReturnValue();

    void modifiedInParam(["java:type:java.util.ArrayList<String>"] StringSeq seq);

    void modifiedOutParam(out ["java:type:java.util.ArrayList<String>"]
StringSeq seq);
};
```


Copyright 2016, ZeroC, Inc.

As you might expect, modifying the mapping for an operation's parameters or return value may require the application to manually convert values from the original mapping to the modified mapping. For example, suppose we want to invoke the `modifiedInParam` operation. The signature of its proxy operation is shown below:

Java

```
void modifiedInParam(java.util.List<String> seq, Ice.Current curr)
```

The metadata changes the mapping of the `seq` parameter to `java.util.List`, which is the default formal type. If a caller has a `StringSeq` value in the original mapping, it must convert the array as shown in the following example:

Java

```
String[] seq = new String[2];
seq[0] = "hi";
seq[1] = "there";
IPrx proxy = ...;
proxy.modifiedInParam(java.util.Arrays.asList(seq));
```

Although we specified the instance type `java.util.ArrayList<String>` for the parameter, we are still able to pass the result of `asList` because its return type (`java.util.List<String>`) is compatible with the parameter's formal type declared by the proxy method. In the case of an operation parameter, the instance type is only relevant to a servant implementation, which may need to make assumptions about the actual type of the parameter.

Mapping for Modified Out Parameters in Java

The mapping for an `out` parameter uses a generated "holder" class to convey the parameter value. If you modify the mapping of an `out` parameter, as discussed in the previous section, it is possible that the holder class for the parameter's unmodified type is no longer compatible with the custom type you have specified. The holder class generated for `StringSeq` is shown below:

Java

```
public final class StringSeqHolder
{
    public
    StringSeqHolder()
    {
    }

    public
    StringSeqHolder(String[] value)
    {
        this.value = value;
    }

    public String[] value;
}
```


Copyright 2016, ZeroC, Inc.

An `out` parameter of type `StringSeq` would normally map to a proxy method that used `StringSeqHolder` to hold the parameter value. When the parameter is modified, as is the case with the `modifiedOutParam` operation, the Slice-to-Java compiler cannot use `StringSeqHolder` to hold an instance of `java.util.List<String>`, because `StringSeqHolder` is only appropriate for the default mapping to a native array.

As a result, the compiler handles these situations using instances of the generic class `Ice.Holder<T>`, where *T* is the parameter's formal type. Consider the following example:

Slice

```
sequence<string> StringSeq;

interface I {
    void modifiedOutParam(out ["java:type:java.util.ArrayList<String>"]
StringSeq seq);
};
```

The compiler generates the following mapping for the `modifiedOutParam` proxy method:

Java

```
void modifiedOutParam(Ice.Holder<java.util.List<java.lang.String> > seq,
Ice.Current curr)
```

The formal type of the parameter is `java.util.List<String>`, therefore the holder class becomes `Ice.Holder<java.util.List<String>>`.

Buffer Types in Java

As of Ice 3.6, you can annotate sequences of certain primitive types with the `java:buffer` metadata tag to change the mapping to use subclasses of `java.nio.Buffer`. This mapping provides several benefits:

- You can pass a buffer to a Slice API instead of creating and filling a temporary array
- If you need to pass a portion of an existing array, you can wrap it with a buffer and avoid an extra copy
- Receiving buffers during a Slice operation also avoids copying by directly referencing the data in Ice's unmarshaling buffer

To use buffers safely, applications must disable caching by setting `Ice.CacheMessageBuffers` to zero.

The following table lists each supported Slice primitive type with its corresponding mapped class:

Primitive	Mapping
`byte`	`java.nio.ByteBuffer`
`short`	`java.nio.ShortBuffer`
`int`	`java.nio.IntBuffer`
`long`	`java.nio.LongBuffer`
`float`	`java.nio.FloatBuffer`
`double`	`java.nio.DoubleBuffer`

Copyright 2016, ZeroC, Inc.

The `java:buffer` tag can be applied to the initial definition of a sequence, in which case the mapping uses the buffer type for all occurrences of that sequence type:

Slice

```
["java:buffer"] sequence<int> Values;

struct Observation {
    int x;
    int y;
    Values measurements;
};
```

We can construct an `Observation` as follows:

Java

```
Observation obs = new Observation();
obs.x = 5;
obs.y = 9;
obs.measurements = java.nio.IntBuffer.allocate(10);
for(int i = 0; i < obs.measurements.capacity(); ++i)
    obs.measurements.put(i, ...);
```

The `java:buffer` tag can also be applied in more limited situations to override a sequence's normal mapping:

Slice

```
sequence<byte> ByteSeq; // Maps to byte[]

struct Page {
    int offset;
    ["java:buffer"] ByteSeq data; // Maps to java.nio.ByteBuffer
};

interface Decoder {
    ["java:buffer"] ByteSeq decode(ByteSeq data);
};
```

In this example, `ByteSeq` maps by default to a `byte` array, but we've overridden the mapping to use a buffer when this type is used as a data member in `Page` and as the return value of the `decode` operation; the input parameter to `decode` uses the default array mapping.

JavaBean Mapping

The Java mapping optionally generates JavaBean-style methods for the data members of class, structure, and exception types.

JavaBean Generated Methods

Copyright 2016, ZeroC, Inc.

For each data member *val* of type *T*, the mapping generates the following methods:

Java

```
public T getVal();
public void setVal(T v);
```

The mapping generates an additional method if *T* is the `bool` type:

Java

```
public boolean isVal();
```

Finally, if *T* is a sequence type with an element type *E*, two methods are generated to provide direct access to elements:

Java

```
public E getVal(int index);
public void setVal(int index, E v);
```

Note that these element methods are only generated for sequence types that use the default mapping.

The Slice-to-Java compiler considers it a fatal error for a JavaBean method of a class data member to conflict with a declared operation of the class. In this situation, you must rename the operation or the data member, or disable the generation of JavaBean methods for the data member in question.

JavaBean Metadata

The JavaBean methods are generated for a data member when the member or its enclosing type is annotated with the `java:getset` meta data. The following example demonstrates both styles of usage:

Copyright 2016, ZeroC, Inc.

Slice

```
sequence<int> IntSeq;

class C {
    ["java:getset"] int i;
    double d;
};

["java:getset"]
struct S {
    bool b;
    string str;
};

["java:getset"]
exception E {
    IntSeq seq;
};
```

JavaBean methods are generated for all members of struct S and exception E, but for only one member of class C. Relevant portions of the generated code are shown below:

Java

```
public class C extends Ice.ObjectImpl
{
    ...

    public int i;

    public int
    getI()
    {
        return i;
    }

    public void
    setI(int _i)
    {
        i = _i;
    }

    public double d;
}

public final class S implements java.lang.Cloneable
{
    public boolean b;
```


Copyright 2016, ZeroC, Inc.

```
    public boolean
    getB()
    {
        return b;
    }

    public void
    setB(boolean _b)
    {
        b = _b;
    }

    public boolean
    isB()
    {
        return b;
    }

    public String str;

    public String
    getStr()
    {
        return str;
    }

    public void
    setStr(String _str)
    {
        str = _str;
    }

    ...
}

public class E extends Ice.UserException
{
    ...

    public int[] seq;

    public int[]
    getSeq()
    {
        return seq;
    }

    public void
    setSeq(int[] _seq)
    {
        seq = _seq;
    }

    public int
    getSeq(int _index)
    {
```


Copyright 2016, ZeroC, Inc.

```
        return seq[_index];
    }

    public void
    setSeq(int _index, int _val)
    {
        seq[_index] = _val;
    }
```


Copyright 2016, ZeroC, Inc.

```
    ...
}
```

Overriding `serialVersionUID`

The Slice-to-Java compiler computes a default value for the `serialVersionUID` member of Slice classes, exceptions and structures. If you prefer, you can override this value using the `java:serialVersionUID` metadata, as shown below:

Slice

```
["java:serialVersionUID:571254925"]
struct Identity
{
    ...
};
```

The specified value will be used in place of the default value in the generated code:

Java

```
public class Identity
{
    ...

    public static final long serialVersionUID = 571254925L;
}
```

By using this metadata, the application assumes responsibility for updating the UID whenever changes to the Slice definition affect the serializable state of the type.

See Also

- Metadata
- Java Mapping for Modules
- Java Mapping for Operations
- Class Inheritance Semantics

Copyright 2016, ZeroC, Inc.

Asynchronous Method Invocation (AMI) in Java

Asynchronous Method Invocation (AMI) is the term used to describe the client-side support for the asynchronous programming model. AMI supports both oneway and twoway requests, but unlike their synchronous counterparts, AMI requests never block the calling thread. When a client issues an AMI request, the Ice run time hands the message off to the local transport buffer or, if the buffer is currently full, queues the request for later delivery. The application can then continue its activities and poll or wait for completion of the invocation, or receive a callback when the invocation completes.

AMI is transparent to the server: there is no way for the server to tell whether a client sent a request synchronously or asynchronously.

On this page:

- Basic Asynchronous API in Java
  - Asynchronous Proxy Methods in Java
  - Asynchronous Exception Semantics in Java
- AsyncResult Interface in Java
- Polling for Completion in Java
- Generic Completion Callbacks in Java
- Sharing State Between begin_ and end_ Methods in Java
- Type-Safe Completion Callbacks in Java
  - Type-Safe Callback Classes in Java
  - Type-Safe Lambda Functions in Java
- Asynchronous Oneway Invocations in Java
- Flow Control in Java
- Asynchronous Batch Requests in Java
- Concurrency Semantics for AMI in Java

Basic Asynchronous API in Java

Consider the following simple Slice definition:

Slice

```
module Demo {
    interface Employees {
        string getName(int number);
    };
};
```

Asynchronous Proxy Methods in Java

Besides the synchronous proxy methods, `slice2java` generates the following asynchronous proxy methods:

Java

```
public interface EmployeesPrx extends Ice.ObjectPrx
{
    // ...

    public Ice.AsyncResult begin_getName(int number);
    public Ice.AsyncResult begin_getName(int number,
java.util.Map<String, String> __ctx);

    public String end_getName(Ice.AsyncResult __result);
}
```


Copyright 2016, ZeroC, Inc.

Four additional overloads of `begin_getName` are generated for use with generic completion callbacks and type-safe completion callbacks.

As you can see, the single `getName` operation results in `begin_getName` and `end_getName` methods. (The `begin_` method is overloaded so you can pass a per-invocation context.)

- The `begin_getName` method sends (or queues) an invocation of `getName`. This method does not block the calling thread.
- The `end_getName` method collects the result of the asynchronous invocation. If, at the time the calling thread calls `end_getName`, the result is not yet available, the calling thread blocks until the invocation completes. Otherwise, if the invocation completed some time before the call to `end_getName`, the method returns immediately with the result.

A client could call these methods as follows:

Java

```
EmployeesPrx e = ...;
Ice.AsyncResult r = e.begin_getName(99);

// Continue to do other things here...

String name = e.end_getName(r);
```

Because `begin_getName` does not block, the calling thread can do other things while the operation is in progress.

Note that `begin_getName` returns a value of type `AsyncResult`. This value contains the state that the Ice run time requires to keep track of the asynchronous invocation. You must pass the `AsyncResult` that is returned by the `begin_` method to the corresponding `end_` method.

The `begin_` method has one parameter for each in-parameter of the corresponding Slice operation. Similarly, the `end_` method has one out-parameter for each out-parameter of the corresponding Slice operation (plus the `AsyncResult` parameter). For example, consider the following operation:

Slice

```
double op(int inp1, string inp2, out bool outp1, out long outp2);
```

The `begin_op` and `end_op` methods have the following signature:

Java

```
Ice.AsyncResult begin_op(int inp1, String inp2);
Ice.AsyncResult begin_op(int inp1, String inp2,
java.util.Map<String, String> __ctx);
double end_op(Ice.BooleanHolder outp1, Ice.LongHolder outp2, Ice.AsyncResult r);
```

Asynchronous Exception Semantics in Java

If an invocation raises an exception, the exception is thrown by the `end_` method, even if the actual error condition for the exception was encountered during the `begin_` method ("on the way out"). The advantage of this behavior is that all exception handling is located with the code that calls the `end_` method (instead of being present twice, once where the `begin_` method is called, and again where the `end_` meth

Copyright 2016, ZeroC, Inc.

od is called).

There is one exception to the above rule: if you destroy the communicator and then make an asynchronous invocation, the `begin_` method throws `CommunicatorDestroyedException`. This is necessary because, once the run time is finalized, it can no longer throw an exception from the `end_` method.

The only other exception that is thrown by the `begin_` and `end_` methods is `java.lang.IllegalArgumentException`. This exception indicates that you have used the API incorrectly. For example, the `begin_` method throws this exception if you call an operation that has a return value or out-parameters on a oneway proxy. Similarly, the `end_` method throws this exception if you use a different proxy to call the `end_` method than the proxy you used to call the `begin_` method, or if the `AsyncResult` you pass to the `end_` method was obtained by calling the `begin_` method for a different operation.

`AsyncResult` Interface in Java

The `AsyncResult` that is returned by the `begin_` method encapsulates the state of the asynchronous invocation:

Java

```
public interface AsyncResult {
    public void cancel();

    public Communicator getCommunicator();
    public Connection getConnection();
    public ObjectPrx getProxy();
    public String getOperation();

    public boolean isCompleted();
    public void waitForCompleted();

    public boolean isSent();
    public void waitForSent();

    public void throwLocalException();

    public boolean sentSynchronously();
}
```

The methods have the following semantics:

- `void cancel()`
  This method prevents a queued invocation from being sent or, if the invocation has already been sent, ignores a reply if the server sends one. `cancel` is a local operation and has no effect on the server. A canceled invocation is considered to be completed, meaning `isCompleted` returns true, and the result of the invocation is an `Ice.InvocationCanceledException`.
- `Communicator getCommunicator()`
  This method returns the communicator that sent the invocation.
- `Connection getConnection()`
  This method returns the connection that was used for the invocation. Note that, for typical asynchronous proxy invocations, this method returns a nil value because the possibility of automatic retries means the connection that is currently in use could change unexpectedly. The `getConnection` method only returns a non-nil value when the `AsyncResult` object is obtained by calling `begin_flushBatchRequests` on a `Connection` object.
- `ObjectPrx getProxy()`
  This method returns the proxy that was used to call the `begin_` method, or nil if the `AsyncResult` object was not obtained via an asynchronous proxy invocation.
- `String getOperation()`

Copyright 2016, ZeroC, Inc.

This method returns the name of the operation.

- `boolean isCompleted()`
  This method returns true if, at the time it is called, the result of an invocation is available, indicating that a call to the `end_` method will not block the caller. Otherwise, if the result is not yet available, the method returns false.
- `void waitForCompleted()`
  This method blocks the caller until the result of an invocation becomes available.
- `boolean isSent()`
  When you call the `begin_` method, the Ice run time attempts to write the corresponding request to the client-side transport. If the transport cannot accept the request, the Ice run time queues the request for later transmission. `isSent` returns true if, at the time it is called, the request has been written to the local transport (whether it was initially queued or not). Otherwise, if the request is still queued or an exception occurred before the request could be sent, `isSent` returns false.
- `void waitForSent()`
  This method blocks the calling thread until a request has been written to the client-side transport, or an exception occurs. After `waitForSent` returns, `isSent` returns true if the request was successfully written to the client-side transport, or false if an exception occurred. In the case of a failure, you can call the corresponding `end_` method or `throwLocalException` to obtain the exception.
- `void throwLocalException()`
  This method throws the local exception that caused the invocation to fail. If no exception has occurred yet, `throwLocalException` does nothing.
- `boolean sentSynchronously()`
  This method returns true if a request was written to the client-side transport without first being queued. If the request was initially queued, `sentSynchronously` returns false (independent of whether the request is still in the queue or has since been written to the client-side transport).

Polling for Completion in Java

The `AsyncResult` methods allow you to poll for call completion. Polling is useful in a variety of cases. As an example, consider the following simple interface to transfer files from client to server:

Slice

```
interface FileTransfer
{
    void send(int offset, ByteSeq bytes);
};
```

The client repeatedly calls `send` to send a chunk of the file, indicating at which offset in the file the chunk belongs. A naïve way to transmit a file would be along the following lines:

Copyright 2016, ZeroC, Inc.

Java

```
FileHandle file = open(...);
FileTransferPrx ft = ...;
int chunkSize = ...;
int offset = 0;
while (!file.eof()) {
    byte[] bs;
    bs = file.read(chunkSize); // Read a chunk
    ft.send(offset, bs);       // Send the chunk
    offset += bs.length;
}
```

This works, but not very well: because the client makes synchronous calls, it writes each chunk on the wire and then waits for the server to receive the data, process it, and return a reply before writing the next chunk. This means that both client and server spend much of their time doing nothing — the client does nothing while the server processes the data, and the server does nothing while it waits for the client to send the next chunk.

Using asynchronous calls, we can improve on this considerably:

Copyright 2016, ZeroC, Inc.

Java

```
FileHandle file = open(...);
FileTransferPrx ft = ...;
int chunkSize = ...;
int offset = 0;

LinkedList<Ice.AsyncResult> results = new LinkedList<Ice.AsyncResult>();
int numRequests = 5;

while (!file.eof()) {
    byte[] bs;
    bs = file.read(chunkSize);

    // Send up to numRequests + 1 chunks asynchronously.
    Ice.AsyncResult r = ft.begin_send(offset, bs);
    offset += bs.length;

    // Wait until this request has been passed to the transport.
    r.waitForSent();
    results.add(r);

    // Once there are more than numRequests, wait for the least
    // recent one to complete.
    while (results.size() > numRequests) {
        Ice.AsyncResult r = results.getFirst();
        results.removeFirst();
        r.waitForCompleted();
    }
}

// Wait for any remaining requests to complete.
while (results.size() > 0) {
    Ice.AsyncResult r = results.getFirst();
    results.removeFirst();
    r.waitForCompleted();
}
```

With this code, the client sends up to `numRequests + 1` chunks before it waits for the least recent one of these requests to complete. In other words, the client sends the next request without waiting for the preceding request to complete, up to the limit set by `numRequests`. In effect, this allows the client to "keep the pipe to the server full of data": the client keeps sending data, so both client and server continuously do work.

Obviously, the correct chunk size and value of `numRequests` depend on the bandwidth of the network as well as the amount of time taken by the server to process each request. However, with a little testing, you can quickly zoom in on the point where making the requests larger or queuing more requests no longer improves performance. With this technique, you can realize the full bandwidth of the link to within a percent or two of the theoretical bandwidth limit of a native socket connection.

Generic Completion Callbacks in Java

The `begin_` method is overloaded to allow you to provide completion callbacks. Here are the corresponding methods for the `getName` operation:

Copyright 2016, ZeroC, Inc.

Java

```
Ice.AsyncResult begin_getName(int number, Ice.Callback __cb);

Ice.AsyncResult begin_getName(int number,
                              java.util.Map<String, String> __ctx,
                              Ice.Callback __cb);
```

The second version of `begin_getName` lets you override the default context. Following the in-parameters, the `begin_` method accepts a parameter of type `Ice.Callback`, which is a callback class with a `completed` method that you must provide. The Ice run time invokes the `completed` method when an asynchronous operation completes. For example:

Java

```
public class MyCallback extends Ice.Callback
{
    public void completed(Ice.AsyncResult r)
    {
        EmployeesPrx e = (EmployeesPrx)r.getProxy();
        try {
            String name = e.end_getName(r);
            System.out.println("Name is: " + name);
        } catch (Ice.LocalException ex) {
            System.err.println("Exception is: " + ex);
        }
    }
}
```

Note that your callback class must derive from `Ice.Callback`. The implementation of your callback method must call the `end_` method. The proxy for the call is available via the `getProxy` method on the `AsyncResult` that is passed by the Ice run time. The return type of `getProxy` is `Ice.ObjectPrx`, so you must down-cast the proxy to its correct type.

Your callback method should catch and handle any exceptions that may be thrown by the `end_` method. If an operation can throw user exceptions, this means that you need an additional catch handler for `Ice.UserException` (or catch all possible user exceptions explicitly). If you allow an exception to escape from the callback method, the Ice run time produces a log entry by default and ignores the exception. (You can disable the log message by setting the property `Ice.Warn.AMICallback` to zero.)

To inform the Ice run time that you want to receive a callback for the completion of the asynchronous call, you pass the callback instance to the `begin_` method:

Java

```
EmployeesPrx e = ...;

MyCallback cb = new MyCallback();
e.begin_getName(99, cb);
```

This is often written using an anonymous class instead:

Copyright 2016, ZeroC, Inc.

Java

```
EmployeesPrx e = ...;

e.begin_getName(
        99,
        new Ice.Callback()
        {
            public void completed(Ice.AsyncResult r)
            {
                EmployeesPrx p = (EmployeesPrx)r.getProxy();
                try {
                    String name = p.end_getName(r);
                    System.out.println("Name is: " + name);
                } catch (Ice.LocalException ex) {
                    System.err.println("Exception: " + ex);
                }
            }
        });
```

An anonymous class is especially useful for callbacks that do only a small amount of work because the code that starts the call and the code that processes the results are physically close together.

Sharing State Between begin_ and end_ Methods in Java

It is common for the `end_` method to require access to some state that is established by the code that calls the `begin_` method. As an example, consider an application that asynchronously starts a number of operations and, as each operation completes, needs to update different user interface elements with the results. In this case, the `begin_` method knows which user interface element should receive the update, and the `end_` method needs access to that element.

Assuming that we have a `Widget` class that designates a particular user interface element, you could pass different widgets by storing the widget to be used as a member of your callback class:

Copyright 2016, ZeroC, Inc.

Java

```
public class MyCallback extends Ice.Callback
{
    public MyCallback(Widget w)
    {
        _w = w;
    }

    private Widget _w;

    public void completed(Ice.AsyncResult r)
    {
        EmployeesPrx e = (EmployeesPrx)r.getProxy();
        try {
            String name = e.end_getName(r);
            _w.writeString(name);
        } catch (Ice.LocalException ex) {
            System.err.println("Exception is: " + ex);
        }
    }
}
```

For this example, we assume that widgets have a `writeString` method that updates the relevant UI element.

When you call the `begin_` method, you pass the appropriate callback instance to inform the `end_` method how to update the display:

Java

```
EmployeesPrx e = ...;
Widget widget1 = ...;
Widget widget2 = ...;

// Invoke the getName operation with different widget callbacks.
e.begin_getName(99, new MyCallback(widget1));
e.begin_getName(24, new MyCallback(widget2));
```

The callback class provides a simple and effective way for you to pass state between the point where an operation is invoked and the point where its results are processed. Moreover, if you have a number of operations that share common state, you can pass the same callback instance to multiple invocations. (If you do this, your callback methods may need to use synchronization.)

Type-Safe Completion Callbacks in Java

The generic callback API is not entirely type-safe:

- You must down-cast the return value of `getProxy` to the correct proxy type before you can call the `end_` method.
- You must call the correct `end_` method to match the operation called by the `begin_` method.
- You must remember to catch exceptions when you call the `end_` method; if you forget to do this, you will not know that the operation failed.

Copyright 2016, ZeroC, Inc.

`slice2java` generates an additional type-safe API that takes care of these chores for you. To use type-safe callbacks, you can either implement a callback class or use lambda functions (in Java 8).

Type-Safe Callback Classes in Java

A callback class must define two callback methods:

- a `response` method that is called if the operation succeeds
- an `exception` method that is called if the operation raises an exception

The class must derive from the base class that is generated by `slice2java`. The name of this base class is `<module>.Callback_<interface>_<operation>`. Here is a callback class for an invocation of the `getName` operation:

Java

```
public class MyCallback extends Demo.Callback_Employees_getName
{
    public void response(String name)
    {
        System.out.println("Name is: " + name);
    }

    public void exception(Ice.LocalException ex)
    {
        System.err.println("Exception is: " + ex);
    }
}
```

The `response` callback parameters depend on the operation signature. If the operation has non-`void` return type, the first parameter of the `response` callback is the return value. The return value (if any) is followed by a parameter for each out-parameter of the corresponding Slice operation, in the order of declaration.

The `exception` callback is invoked if the invocation fails because of an Ice run time exception. If the Slice operation can also raise user exceptions, your callback class must supply an additional overloading of `exception` that accepts an argument of type `Ice.UserException`.

The proxy methods are overloaded to accept this callback instance:

Java

```
Ice.AsyncResult begin_getName(int number,
                              Callback_Employees_getName __cb);

Ice.AsyncResult begin_getName(int number,
                              java.util.Map<String, String> __ctx,
                              Callback_Employees_getName __cb);
```

You pass the callback to an invocation as you would with the generic API:

Copyright 2016, ZeroC, Inc.

Java

```
EmployeesPrx e = ...;

MyCallback cb = new MyCallback();
e.begin_getName(99, cb);
```

Type-Safe Lambda Functions in Java

If you're using Java 8, you can implement your type-safe callbacks using in-line lambda functions. The proxy methods are overloaded to accept these functions:

Java

```
Ice.AsyncResult begin_getName(int number,
                              IceInternal.Functional_GenericCallback1<String>
__responseCb,

IceInternal.Functional_GenericCallback1<Ice.Exception> __exceptionCb);

Ice.AsyncResult begin_getName(int number,
                              java.util.Map<String, String> __ctx,
                              IceInternal.Functional_GenericCallback1<String>
__responseCb,

IceInternal.Functional_GenericCallback1<Ice.Exception> __exceptionCb);
```

The names of the internal interfaces used in the API are not important; what matters are their parameterized types, which tell you the arguments that your lambda functions must accept. For example, we can call `begin_getName` as follows:

Java

```
EmployeesPrx e = ...;

e.begin_getName(99,
    (String name) -> {
        System.out.println("Name is: " + name);
    },
    (Ice.Exception ex) -> {
        System.err.println("Exception is: " + ex);
    });
```

Asynchronous Oneway Invocations in Java

Copyright 2016, ZeroC, Inc.

You can invoke operations via oneway proxies asynchronously, provided the operation has `void` return type, does not have any out-parameters, and does not raise user exceptions. If you call the `begin_` method on a oneway proxy for an operation that returns values or raises a user exception, the `begin_` method throws an `IllegalArgumentException`.

The callback methods looks exactly as for a twoway invocation. For the generic API, the Ice run time does not call the `completed` callback method unless the invocation raised an exception during the `begin_` method ("on the way out"). For the type-safe API, the `response` meth od is never called.

Flow Control in Java

Asynchronous method invocations never block the thread that calls the `begin_` method: the Ice run time checks to see whether it can write the request to the local transport. If it can, it does so immediately in the caller's thread. (In that case, `AsyncResult.sentSynchronously` returns true.) Alternatively, if the local transport does not have sufficient buffer space to accept the request, the Ice run time queues the request internally for later transmission in the background. (In that case, `AsyncResult.sentSynchronously` returns false.)

This creates a potential problem: if a client sends many asynchronous requests at the time the server is too busy to keep up with them, the requests pile up in the client-side run time until, eventually, the client runs out of memory.

The API provides a way for you to implement flow control by counting the number of requests that are queued so, if that number exceeds some threshold, the client stops invoking more operations until some of the queued operations have drained out of the local transport.

For the generic API, you can override the `sent` method:

Java

```
public class MyCallback extends Ice.Callback
{
    public void completed(Ice.AsyncResult r)
    {
        // ...
    }

    public void sent(Ice.AsyncResult r)
    {
        // ...
    }
}
```

You inform the Ice run time that you want to be informed when a call has been passed to the local transport as usual:

Java

```
e.begin_getName(99, new MyCallback());
```

If the Ice run time can immediately pass the request to the local transport, it does so and invokes the `sent` method from the thread that calls the `begin_` method. On the other hand, if the run time has to queue the request, it calls the `sent` method from a different thread once it has written the request to the local transport. In addition, you can find out from the `AsyncResult` that is returned by the `begin_` method whether the request was sent synchronously or was queued, by calling `sentSynchronously`.

For the generic API, the `sent` method has the following signature:

Copyright 2016, ZeroC, Inc.

Java

```
void sent(Ice.AsyncResult r);
```

For the type-safe API, the signature is:

Java

```
void sent(boolean sentSynchronously);
```

For the generic API, you can find out whether the request was sent synchronously by calling `sentSynchronously` on the `AsyncResult`. For the type-safe API, the boolean `sentSynchronously` parameter provides the same information.

The `sent` methods allow you to limit the number of queued requests by counting the number of requests that are queued and decrementing the count when the Ice run time passes a request to the local transport.

Asynchronous Batch Requests in Java

Applications that send batched requests can either flush a batch explicitly or allow the Ice run time to flush automatically. The proxy method `ice_flushBatchRequests` performs an immediate flush using the synchronous invocation model and may block the calling thread until the entire message can be sent. Ice also provides asynchronous versions of this method so you can flush batch requests asynchronously.

`begin_ice_flushBatchRequests` and `end_ice_flushBatchRequests` are proxy methods that flush any batch requests queued by that proxy.

In addition, similar methods are available on the communicator and the `Connection` object that is returned by `AsyncResult.getConnection`. These methods flush batch requests sent via the same communicator and via the same connection, respectively.

Concurrency Semantics for AMI in Java

The Ice run time always invokes your callback methods from a separate thread, with one exception: it calls the `sent` callback from the thread calling the `begin_` method if the request could be sent synchronously. In the `sent` callback, you know which thread is calling the callback by looking at the `sentSynchronously` member or parameter.

See Also

- Request Contexts
- Batched Invocations
- Collocated Invocation and Dispatch

Copyright 2016, ZeroC, Inc.

Using the Slice Compiler for Java

Redirection Notice
This page will redirect to slice2java Command-Line Options in about one second.

Copyright 2016, ZeroC, Inc.

slice2java Command-Line Options

The Slice-to-Java compiler, `slice2java`, offers the following command-line options in addition to the standard options:

- `--tie`
  Generate tie classes.
- `--impl`
  Generate sample implementation files. This option will not overwrite an existing file.
- `--impl-tie`
  Generate sample implementation files using tie classes. This option will not overwrite an existing file.
- `--checksum` *`CLASS`*
  Generate checksums for Slice definitions into the class *`CLASS`*. The given class name may optionally contain a package specifier. The generated class contains checksums for all of the Slice files being translated by this invocation of the compiler. For example, the command below causes `slice2java` to generate the file `Checksums.java` containing the checksums for the Slice definitions in `File1.ice` and `File2.ice`:

  ```
  slice2java --checksum Checksums File1.ice File2.ice
  ```

- `--stream`
  Generate streaming helper methods for Slice types.
- `--meta` *`META`*
  Define the global metadata directive *`META`*. Using this option is equivalent to defining the global metadata *`META`* in each named Slice file, as well as in any file included by a named Slice file. Global metadata specified with `--meta` overrides any corresponding global metadata directive in the files being compiled.
- `--list-generated`
  Emit a list of generated files in XML format.

See Also

- Using the Slice Compilers
- Using Slice Checksums in Java
- Tie Classes in Java
- Streaming Interfaces

Copyright 2016, ZeroC, Inc.

Using Slice Checksums in Java

The Slice compilers can optionally generate checksums of Slice definitions. For `slice2java`, the `--checksum` option causes the compiler to generate a new Java class that adds checksums to a static map member. Assuming we supplied the option `--checksum Checksums` to `slice2java`, the generated class `Checksums.java` looks like this:

Java

```
public class Checksums {
    public static java.util.Map<String, String> checksums;
}
```

The read-only map `checksums` is initialized automatically prior to first use; no action is required by the application.

In order to verify a server's checksums, a client could simply compare the dictionaries using the `equals` method. However, this is not feasible if it is possible that the server might return a superset of the client's checksums. A more general solution is to iterate over the local checksums as demonstrated below:

Java

```
java.util.Map<String, String> serverChecksums = ...
java.util.Iterator<java.util.Map.Entry<String, String>> i =
    Checksums.checksums.entrySet().iterator();
while(i.hasNext()) {
    java.util.Map.Entry<String, String> e = i.next();
    String id = e.getKey();
    String checksum = e.getValue();
    String serverChecksum = serverChecksums.get(id);
    if (serverChecksum == null) {
        // No match found for type id!
    } else if (!checksum.equals(serverChecksum)) {
        // Checksum mismatch!
    }
}
```

In this example, the client first verifies that the server's dictionary contains an entry for each Slice type ID, and then it proceeds to compare the checksums.

See Also

- Slice Checksums

Copyright 2016, ZeroC, Inc.

Example of a File System Client in Java

This page presents the source code for a very simple client to access a server that implements the file system we developed in Slice for a Simple File System. The Java code hardly differs from the code you would write for an ordinary Java program. This is one of the biggest advantages of using Ice: accessing a remote object is as easy as accessing an ordinary, local Java object. This allows you to put your effort where you should, namely, into developing your application logic instead of having to struggle with arcane networking APIs. This is true for the server side as well, meaning that you can develop distributed applications easily and efficiently.

We now have seen enough of the client-side Java mapping to develop a complete client to access our remote file system. For reference, here is the Slice definition once more:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };

    exception GenericError {
        string reason;
    };

    sequence<string> Lines;

    interface File extends Node {
        idempotent Lines read();
        idempotent void write(Lines text) throws GenericError;
    };

    sequence<Node*> NodeSeq;

    interface Directory extends Node {
        idempotent NodeSeq list();
    };
};
```

To exercise the file system, the client does a recursive listing of the file system, starting at the root directory. For each node in the file system, the client shows the name of the node and whether that node is a file or directory. If the node is a file, the client retrieves the contents of the file and prints them.

The body of the client code looks as follows:

Java

```
import Filesystem.*;

public class Client {

    // Recursively print the contents of directory "dir" in
    // tree fashion.  For files, show the contents of each file.
    // The "depth" parameter is the current nesting level
    // (for indentation).

    static void
```


Copyright 2016, ZeroC, Inc.

```
    listRecursive(DirectoryPrx dir, int depth)
    {
        char[] indentCh = new char[++depth];
        java.util.Arrays.fill(indentCh, '\t');
        String indent = new String(indentCh);

        NodePrx[] contents = dir.list();

        for (int i = 0; i < contents.length; ++i) {
            DirectoryPrx subdir = DirectoryPrxHelper.checkedCast(contents[i]);
            FilePrx file = FilePrxHelper.uncheckedCast(contents[i]);
            System.out.println(indent + contents[i].name() +
                (subdir != null ? " (directory):" : " (file):"));
            if (subdir != null) {
                listRecursive(subdir, depth);
            } else {
                String[] text = file.read();
                for (int j = 0; j < text.length; ++j)
                    System.out.println(indent + "\t" + text[j]);
            }
        }
    }

    public static void
    main(String[] args)
    {
        int status = 0;
        Ice.Communicator ic = null;
        try {
            // Create a communicator
            //
            ic = Ice.Util.initialize(args);

            // Create a proxy for the root directory
            //
            Ice.ObjectPrx base = ic.stringToProxy("RootDir:default -p 10000");
            if (base == null)
                throw new RuntimeException("Cannot create proxy");

            // Down-cast the proxy to a Directory proxy
            //
            DirectoryPrx rootDir = DirectoryPrxHelper.checkedCast(base);
            if (rootDir == null)
                throw new RuntimeException("Invalid proxy");

            // Recursively list the contents of the root directory
            //
            System.out.println("Contents of root directory:");
            listRecursive(rootDir, 0);
        } catch (Ice.LocalException e) {
            e.printStackTrace();
            status = 1;
        } catch (Exception e) {
            System.err.println(e.getMessage());
            status = 1;
        }
```


Copyright 2016, ZeroC, Inc.

```
        if (ic != null) {
            // Clean up
            //
            try {
                ic.destroy();
            } catch (Exception e) {
                System.err.println(e.getMessage());
                status = 1;
            }
        }
        System.exit(status);
```


Copyright 2016, ZeroC, Inc.

```
    }
}
```

After importing the `Filesystem` package, the `Client` class defines two methods: `listRecursive`, which is a helper function to print the contents of the file system, and `main`, which is the main program. Let us look at `main` first:

1. The structure of the code in `main` follows what we saw in Hello World Application. After initializing the run time, the client creates a proxy to the root directory of the file system. For this example, we assume that the server runs on the local host and listens using the default protocol (TCP/IP) at port 10000. The object identity of the root directory is known to be `RootDir`.
2. The client down-casts the proxy to `DirectoryPrx` and passes that proxy to `listRecursive`, which prints the contents of the file system.

Most of the work happens in `listRecursive`. The function is passed a proxy to a directory to list, and an indent level. (The indent level increments with each recursive call and allows the code to print the name of each node at an indent level that corresponds to the depth of the tree at that node.) `listRecursive` calls the `list` operation on the directory and iterates over the returned sequence of nodes:

1. The code does a `checkedCast` to narrow the `Node` proxy to a `Directory` proxy, as well as an `uncheckedCast` to narrow the `Node` proxy to a `File` proxy. Exactly one of those casts will succeed, so there is no need to call `checkedCast` twice: if the `Node` *is-a* `Directory`, the code uses the `DirectoryPrx` returned by the `checkedCast`; if the `checkedCast` fails, we *know* that the Node *is-a* File and, therefore, an `uncheckedCast` is sufficient to get a `FilePrx`.
   In general, if you know that a down-cast to a specific type will succeed, it is preferable to use an `uncheckedCast` instead of a `checkedCast` because an `uncheckedCast` does not incur any network traffic.
2. The code prints the name of the file or directory and then, depending on which cast succeeded, prints `"(directory)"` or `"(file)"` following the name.
3. The code checks the type of the node:
   - If it is a directory, the code recurses, incrementing the indent level.
   - If it is a file, the code calls the `read` operation on the file to retrieve the file contents and then iterates over the returned sequence of lines, printing each line.

Assume that we have a small file system consisting of two files and a directory as follows:

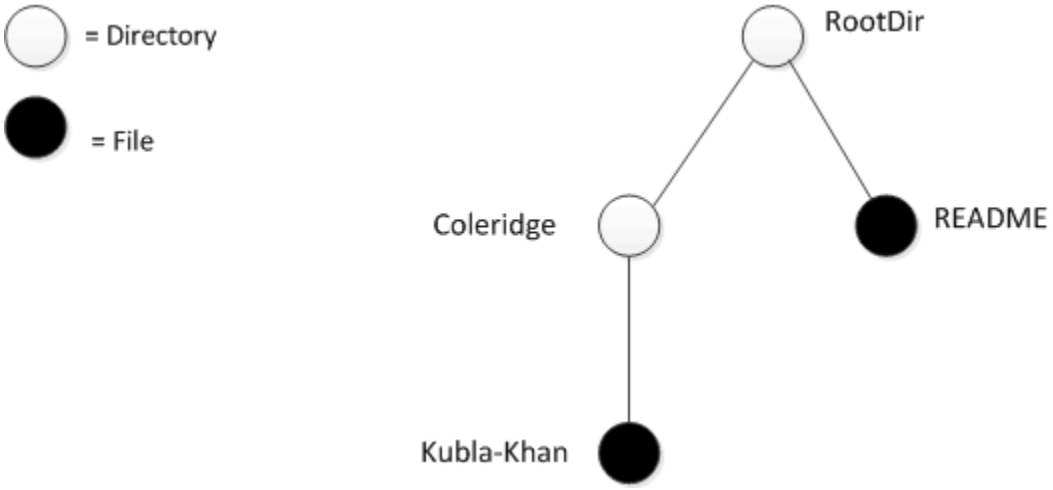

A small file system.

The output produced by the client for this file system is:

Copyright 2016, ZeroC, Inc.

```
Contents of root directory:
        README (file):
                This file system contains a collection of poetry.
        Coleridge (directory):
                Kubla_Khan (file):
                        In Xanadu did Kubla Khan
                        A stately pleasure-dome decree:
                        Where Alph, the sacred river, ran
                        Through caverns measureless to man
                        Down to a sunless sea.
```

Note that, so far, our client (and server) are not very sophisticated:

- The protocol and address information are hard-wired into the code.
- The client makes more remote procedure calls than strictly necessary; with minor redesign of the Slice definitions, many of these calls can be avoided.

We will see how to address these shortcomings in our discussions of IceGrid and object life cycle.

See Also

- Hello World Application
- Slice for a Simple File System
- Example of a File System Server in Java
- Object Life Cycle
- IceGrid

Copyright 2016, ZeroC, Inc.

Server-Side Slice-to-Java Mapping

The mapping for Slice data types to Java is identical on the client side and server side. This means that everything in Client-Side Slice-to-Java Mapping also applies to the server side. However, for the server side, there are a few additional things you need to know — specifically how to:

- Initialize and finalize the server-side run time
- Implement servants
- Pass parameters and throw exceptions
- Create servants and register them with the Ice run time

Because the mapping for Slice data types is identical for clients and servers, the server-side mapping only adds a few additional mechanisms to the client side: a small API to initialize and finalize the run time, plus a few rules for how to derive servant classes from skeletons and how to register servants with the server-side run time.

Although the examples we present are very simple, they accurately reflect the basics of writing an Ice server. Of course, for more sophisticated servers, you will be using additional APIs, for example, to improve performance or scalability. However, these APIs are all described in Slice, so, to use these APIs, you need not learn any Java mapping rules beyond those we described here.

Topics

- The Server-Side main Method in Java
- Server-Side Java Mapping for Interfaces
- Parameter Passing in Java
- Raising Exceptions in Java
- Tie Classes in Java
- Object Incarnation in Java
- Asynchronous Method Dispatch (AMD) in Java
- Example of a File System Server in Java

Copyright 2016, ZeroC, Inc.

The Server-Side main Method in Java

On this page:

- A Basic main Method in Java
- The Ice.Application Class in Java
  - Using Ice.Application on the Client Side in Java
  - Catching Signals in Java
  - Ice.Application and Properties in Java
  - Limitations of Ice.Application in Java

A Basic `main` Method in Java

The main entry point to the Ice run time is represented by the local Slice interface `Ice::Communicator`. As for the client side, you must initialize the Ice run time by calling `Ice.Util.initialize` before you can do anything else in your server. `Ice.Util.initialize` returns a reference to an instance of an `Ice.Communicator`:

Java

```
public class Server {
    public static void
    main(String[] args)
    {
        int status = 0;
        Ice.Communicator ic = null;
        try {
            ic = Ice.Util.initialize(args);
            // ...
        } catch (Exception e) {
            e.printStackTrace();
            status = 1;
        }
        // ...
    }
}
```

`Ice.Util.initialize` accepts the argument vector that is passed to `main` by the operating system. The function scans the argument vector for any command-line options that are relevant to the Ice run time, but does not remove those options. If anything goes wrong during initialization, `initialize` throws an exception.

The semantics of Java arrays prevents `Ice.Util.initialize` from modifying the size of the argument vector. However, another overloading of `Ice.Util.initialize` is provided that allows the application to obtain a new argument vector with the Ice options removed.

Before leaving your `main` function, you *must* call `Communicator.destroy`. The `destroy` operation is responsible for finalizing the Ice run time. In particular, `destroy` waits for any operation implementations that are still executing in the server to complete. In addition, `destroy` ensures that any outstanding threads are joined with and reclaims a number of operating system resources, such as file descriptors and memory. Never allow your `main` function to terminate without calling `destroy` first; doing so has undefined behavior.

The general shape of our server-side `main` function is therefore as follows:

Copyright 2016, ZeroC, Inc.

Java

```
public class Server {
    public static void
    main(String[] args)
    {
        int status = 0;
        Ice.Communicator ic = null;
        try {
            ic = Ice.Util.initialize(args);
            // ...
        } catch (Exception e) {
            e.printStackTrace();
            status = 1;
        }
        if (ic != null) {
            try {
                ic.destroy();
            } catch (Exception e) {
                e.printStackTrace();
                status = 1;
            }
        }
        System.exit(status);
    }
}
```

Note that the code places the call to `Ice.Util.initialize` into a `try` block and takes care to return the correct exit status to the operating system. Also note that an attempt to destroy the communicator is made only if the initialization succeeded.

The `Ice.Application` Class in Java

The preceding structure for the `main` function is so common that Ice offers a class, `Ice.Application`, that encapsulates all the correct initialization and finalization activities. The synopsis of the class is as follows (with some detail omitted for now):

Copyright 2016, ZeroC, Inc.

Java

```
package Ice;

public enum SignalPolicy { HandleSignals, NoSignalHandling }

public abstract class Application {
    public Application()

    public Application(SignalPolicy signalPolicy)

    public final int main(String appName, String[] args)

    public final int main(String appName, String[] args, String configFile)

    public final int main(String appName, String[] args,
InitializationData initData)

    public abstract int run(String[] args)

    public static String appName()

    public static Communicator communicator()

    // ...
}
```

The intent of this class is that you specialize `Ice.Application` and implement the abstract `run` method in your derived class. Whatever code you would normally place in `main` goes into the `run` method instead. Using `Ice.Application`, our program looks as follows:

Java

```
public class Server extends Ice.Application {
    public int
    run(String[] args)
    {
        // Server code here...

        return 0;
    }

    public static void
    main(String[] args)
    {
        Server app = new Server();
        int status = app.main("Server", args);
        System.exit(status);
    }
}
```


Copyright 2016, ZeroC, Inc.

Note that `Application.main` is overloaded: you can pass an optional file name or an `InitializationData` structure.

If you pass a configuration file name to `main`, the property settings in this file are overridden by settings in a file identified by the `ICE_CONFIG` environment variable (if defined). Property settings supplied on the command line take precedence over all other settings.

The `Application.main` function does the following:

1. It installs an exception handler for `java.lang.Exception`. If your code fails to handle an exception, `Application.main` prints the name of an exception and a stack trace on `System.err` before returning with a non-zero return value.
2. It initializes (by calling `Ice.Util.initialize`) and finalizes (by calling `Communicator.destroy`) a communicator. You can get access to the communicator for your server by calling the static `communicator` accessor.
3. It scans the argument vector for options that are relevant to the Ice run time and removes any such options. The argument vector that is passed to your `run` method therefore is free of Ice-related options and only contains options and arguments that are specific to your application.
4. It provides the name of your application via the static `appName` member function. The return value from this call is the first argument in the call to `Application.main`, so you can get at this name from anywhere in your code by calling `Ice.Application.appName` (which is usually required for error messages). In the example above, the return value from `appName` would be `Server`.
5. It installs a shutdown hook that properly shuts down the communicator.
6. It installs a per-process logger if the application has not already configured one. The per-process logger uses the value of the `Ice.ProgramName` property as a prefix for its messages and sends its output to the standard error channel. An application can also specify an alternate logger.

Using `Ice.Application` ensures that your program properly finalizes the Ice run time, whether your server terminates normally or in response to an exception. We recommend that all your programs use this class; doing so makes your life easier. In addition, `Ice.Application` also provides features for signal handling and configuration that you do not have to implement yourself when you use this class.

Using `Ice.Application` on the Client Side in Java

You can use `Ice.Application` for your clients as well: simply implement a class that derives from `Ice.Application` and place the client code into its `run` method. The advantage of this approach is the same as for the server side: `Ice.Application` ensures that the communicator is destroyed correctly even in the presence of exceptions.

Catching Signals in Java

The simple server we developed in Hello World Application had no way to shut down cleanly: we simply interrupted the server from the command line to force it to exit. Terminating a server in this fashion is unacceptable for many real-life server applications: typically, the server has to perform some cleanup work before terminating, such as flushing database buffers or closing network connections. This is particularly important on receipt of a signal or keyboard interrupt to prevent possible corruption of database files or other persistent data.

Java does not provide direct support for signals, but it does allow an application to register a *shutdown hook* that is invoked when the JVM is shutting down. There are several events that trigger JVM shutdown, such as a call to `System.exit` or an interrupt signal from the operating system, but the shutdown hook is not provided with the reason for the shut down.

`Ice.Application` registers a shutdown hook by default, allowing you to cleanly terminate your application prior to JVM shutdown.

Java

```
package Ice;

public abstract class Application {
    // ...

    synchronized public static void destroyOnInterrupt()
    synchronized public static void shutdownOnInterrupt()
    synchronized public static void setInterruptHook(Thread t)
    synchronized public static void defaultInterrupt()
    synchronized public static boolean interrupted()
}
```

The functions behave as follows:

Copyright 2016, ZeroC, Inc.

- `destroyOnInterrupt`
  This function installs a shutdown hook that calls `destroy` on the communicator. This is the default behavior.
- `shutdownOnInterrupt`
  This function installs a shutdown hook that calls `shutdown` on the communicator.
- `setInterruptHook`
  This function installs a custom shutdown hook that takes responsibility for performing whatever action is necessary to terminate the application. Refer to the Java documentation for `Runtime.addShutdownHook` for more information on the semantics of shutdown hooks.
- `defaultInterrupt`
  This function removes the shutdown hook.
- `interrupted`
  This function returns true if the shutdown hook caused the communicator to shut down, false otherwise. This allows us to distinguish intentional shutdown from a forced shutdown that was caused by the JVM. This is useful, for example, for logging purposes.

By default, `Ice.Application` behaves as if `destroyOnInterrupt` was invoked, therefore our server `main` function requires no change to ensure that the program terminates cleanly on JVM shutdown. (You can disable this default shutdown hook by passing the enumerator `NoSignalHandling` to the constructor. In that case, shutdown is not intercepted and terminates the VM.) However, we add a diagnostic to report the occurrence, so our `main` function now looks like:

Java

```
public class Server extends Ice.Application {
    public int
    run(String[] args)
    {
        // Server code here...

        if (interrupted())
            System.err.println(appName() + ": terminating");

        return 0;
    }

    public static void
    main(String[] args)
    {
        Server app = new Server();
        int status = app.main("Server", args);
        System.exit(status);
    }
}
```

During the course of normal execution, the JVM does not terminate until all non-daemon threads have completed. If an interrupt occurs, the JVM ignores the status of active threads and terminates as soon as it has finished invoking all of the installed shutdown hooks.

In a subclass of `Ice.Application`, the default shutdown hook (as installed by `destroyOnInterrupt`) blocks until the application's main thread completes. As a result, an interrupted application may not terminate successfully if the main thread is blocked. For example, this can occur in an interactive application when the main thread is waiting for console input. To remedy this situation, the application can install an alternate shutdown hook that does not wait for the main thread to finish:

Copyright 2016, ZeroC, Inc.

Java

```
public class Server extends Ice.Application {
    class ShutdownHook extends Thread {
        public void
        run()
        {
            try
            {
                communicator().destroy();
            }
            catch(Ice.LocalException ex)
            {
                ex.printStackTrace();
            }
        }
    }

    public int
    run(String[] args)
    {
        setInterruptHook(new ShutdownHook());

        // ...
    }
}
```

After replacing the default shutdown hook using `setInterruptHook`, the JVM will terminate as soon as the communicator is destroyed.

`Ice.Application` and Properties in Java

Apart from the functionality shown in this section, `Ice.Application` also takes care of initializing the Ice run time with property values. Properties allow you to configure the run time in various ways. For example, you can use properties to control things such as the thread pool size or port number for a server. The `main` function of `Ice.Application` is overloaded; the second version allows you to specify the name of a configuration file that will be processed during initialization.

Limitations of `Ice.Application` in Java

`Ice.Application` is a singleton class that creates a single communicator. If you are using multiple communicators, you cannot use `Ice.Application`. Instead, you must structure your code as we saw in Hello World Application (taking care to always destroy the communicator).
See Also

- Hello World Application
- Properties and Configuration
- Communicator Initialization
- Logger Facility

Copyright 2016, ZeroC, Inc.

Server-Side Java Mapping for Interfaces

The server-side mapping for interfaces provides an up-call API for the Ice run time: by implementing member functions in a servant class, you provide the hook that gets the thread of control from the Ice server-side run time into your application code.

On this page:

- Skeleton Classes in Java
- Servant Classes in Java
  - Normal and idempotent Operations in Java

Skeleton Classes in Java

On the client side, interfaces map to proxy classes. On the server side, interfaces map to *skeleton* classes. A skeleton is a class that has a pure virtual member function for each operation on the corresponding interface. For example, consider our Slice definition for the `Node` interface:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };
    // ...
};
```

The Slice compiler generates the following definition for this interface:

Copyright 2016, ZeroC, Inc.

Java

```
package Filesystem;

public interface _NodeOperations
{
    String name(Ice.Current current);
}

public interface _NodeOperationsNC
{
    String name();
}

public interface Node extends Ice.Object,
                              _NodeOperations,
                              _NodeOperationsNC
{
    public static final String ice_staticId = "::Filesystem::Node";
    public static final long serialVersionUID = ...;
}

public abstract class _NodeDisp extends Ice.ObjectImpl
                                implements Node
{
    // Mapping-internal code here...
}
```

The important points to note here are:

- As for the client side, Slice modules are mapped to Java packages with the same name, so the skeleton class definitions are part of the `Filesystem` package.
- For each Slice interface <*`interface-name`*>, the compiler generates Java interfaces `_`<*`interface-name`*>`Operations` and `_`<*`interface-name`*>`OperationsNC` (`_NodeOperations` and `_NodeOperationsNC` in this example). These interfaces contain a method for each operation in the Slice interface. (You can ignore the `Ice.Current` parameter for now.)
- For each Slice interface <*`interface-name`*>, the compiler generates a Java interface <*`interface-name`*> (`Node` in this example). That interface extends `Ice.Object` as well as the two operations interfaces, and defines the constant `ice_staticId` with the corresponding Slice type ID.
- For each Slice interface <*`interface-name`*>, the compiler generates an abstract class `_`<*`interface-name`*>`Disp` (`_NodeDisp` in this example). This abstract class is the actual skeleton class; it is the base class from which you derive your servant class.

Servant Classes in Java

In order to provide an implementation for an Ice object, you must create a servant class that inherits from the corresponding skeleton class. For example, to create a servant for the `Node` interface, you could write:

Copyright 2016, ZeroC, Inc.

Java

```
package Filesystem;

public final class NodeI extends _NodeDisp {

    public NodeI(String name)
    {
        _name = name;
    }

    public String name(Ice.Current current)
    {
        return _name;
    }

    private String _name;
}
```

By convention, servant classes have the name of their interface with an `I`-suffix, so the servant class for the `Node` interface is called `NodeI`. (This is a convention only: as far as the Ice run time is concerned, you can choose any name you prefer for your servant classes.) Note that `NodeI` extends `_NodeDisp`, that is, it derives from its skeleton class.

As far as Ice is concerned, the `NodeI` class must implement only a single method: the `name` method that it inherits from its skeleton. This makes the servant class a concrete class that can be instantiated. You can add other member functions and data members as you see fit to support your implementation. For example, in the preceding definition, we added a `_name` member and a constructor. (Obviously, the constructor initializes the `_name` member and the `name` function returns its value.)

Normal and `idempotent` Operations in Java

Whether an operation is an ordinary operation or an `idempotent` operation has no influence on the way the operation is mapped. To illustrate this, consider the following interface:

Slice

```
interface Example {
                void   normalOp();
    idempotent void   idempotentOp();
    idempotent string readonlyOp();
};
```

The operations class for this interface looks like this:

Copyright 2016, ZeroC, Inc.

Java

```
public interface _ExampleOperations
{
    void normalOp(Ice.Current current);
    void idempotentOp(Ice.Current current);
    String readonlyOp(Ice.Current current);
}
```

Note that the signatures of the member functions are unaffected by the `idempotent` qualifier.

See Also

- Slice for a Simple File System
- Java Mapping for Interfaces
- Parameter Passing in Java
- Raising Exceptions in Java
- Tie Classes in Java
- The Current Object

Copyright 2016, ZeroC, Inc.

Parameter Passing in Java

For each parameter of a Slice operation, the Java mapping generates a corresponding parameter for the method in the _<*interface-name*>Operations interface. In addition, every operation has a trailing parameter of type Ice.Current. For example, the name operation of the Node interface has no parameters, but the name method of the _NodeOperations interface has a single parameter of type Ice.Current. We will ignore this parameter for now.

Passing Required Parameters in Java

Parameter passing on the server side follows the rules for the client side. To illustrate the rules, consider the following interface that passes string parameters in all possible directions:

Slice

```
module M {
    interface Example {
        string op(string sin, out string sout);
    };
};
```

The generated skeleton class for this interface looks as follows:

Java

```
public interface _ExampleOperations
{
    String op(String sin, Ice.StringHolder sout, Ice.Current current);
}
```

As you can see, there are no surprises here. For example, we could implement op as follows:

Java

```
public final class ExampleI extends M._ExampleDisp {

    public String op(String sin, Ice.StringHolder sout, Ice.Current current)
    {
        System.out.println(sin);     // In params are initialized
        sout.value = "Hello World!"; // Assign out param
        return "Done";
    }
}
```

This code is in no way different from what you would normally write if you were to pass strings to and from a function; the fact that remote procedure calls are involved does not impact on your code in any way. The same is true for parameters of other types, such as proxies, classes, or dictionaries: the parameter passing conventions follow normal Java rules and do not require special-purpose API calls.

Passing Optional Parameters in Java

Suppose we modify the example above to use optional parameters:

Copyright 2016, ZeroC, Inc.

Slice

```
module M {
    interface Example {
        optional(1) string op(optional(2) string sin, out optional(3) string sout);
    };
};
```

The generated skeleton now looks like this:

Java

```
public interface _ExampleOperations
{
    String op(Ice.Optional<String> sin, Ice.StringHolder sout, Ice.Current
current);
}
```

The default mapping treats optional out parameters and return values as if they are required. If your servant needs the ability to return an optional value, you must add the `java:optional` metadata tag:

Slice

```
module M {
    interface Example {
        ["java:optional"]
        optional(1) string op(optional(2) string sin, out optional(3) string sout);
    };
};
```

This tag can be applied to an interface if you want to use the optional mapping for all of the operations in that interface, or to individual operations as shown here. With this change, the mapping now returns optional values:

Java

```
public interface _ExampleOperations
{
    Ice.Optional<String> op(Ice.Optional<String> sin, Ice.Optional<String> sout,
Ice.Current current);
}
```

The `java:optional` tag affects the return value and all out parameters; it is not possible to modify the mapping only for certain parameters.

See Also

- Server-Side Java Mapping for Interfaces
- Java Mapping for Operations
- Raising Exceptions in Java
- Tie Classes in Java

Copyright 2016, ZeroC, Inc.

- The Current Object

Copyright 2016, ZeroC, Inc.

Raising Exceptions in Java

Servant Exceptions in Java

To throw an exception from an operation implementation, you simply instantiate the exception, initialize it, and throw it. For example:

Java

```
// ...

public void
write(String[] text, Ice.Current current)
    throws GenericError
{
    try
    {
        // Try to write file contents here...
    }
    catch(Exception ex)
    {
        throw new GenericError("Exception during write operation", ex);
    }
}
```

Note that, for this example, we have supplied the optional second parameter to the `GenericError` constructor. This parameter sets the inner exception and preserves the original cause of the error for later diagnosis.

If you throw an arbitrary Java run-time exception (such as a `ClassCastException`), the Ice run time catches the exception and then returns an `UnknownException` to the client. Similarly, if you throw an "impossible" user exception (a user exception that is not listed in the exception specification of the operation), the client receives an `UnknownUserException`.

If you throw an Ice run-time exception, such as `MemoryLimitException`, the client receives an `UnknownLocalException`. For that reason, you should never throw system exceptions from operation implementations. If you do, all the client will see is an `UnknownLocalException`, which does not tell the client anything useful.

Three run-time exceptions are treated specially and not changed to `UnknownLocalException` when returned to the client: `ObjectNotExistException`, `OperationNotExistException`, and `FacetNotExistException`.

JVM Error Semantics

Servant implementations might inadvertently raise low-level errors during the course of their operation. These exceptions, which are subclasses of `java.lang.Error`, should not normally be trapped by the servant because they often indicate the occurrence of a serious, unrecoverable situation. For example, the JVM throws `StackOverflowError` when a recursive method has used all available stack space.

The Ice run time traps instances of `java.lang.Error` thrown by servants and then attempts to log the exception and send an `UnknownException` to the client. This attempt may or may not succeed, depending on the nature of the error and the condition of the JVM. As an example, the servant implementation might raise `OutOfMemoryError` and Ice's attempt to log the error and send a response could also fail due to lack of memory.

For an occurrence of `OutOfMemoryError` or `AssertionError`, Ice does not re-throw the error after logging a message and sending a response. For all other subclasses of `Error`, Ice re-throws the error so that the JVM's normal error-handling strategy will execute.

When the JVM raises a subclass of `Error`, it usually means that a significant problem has occurred. Ice tries to send an `UnknownException` to the client (for twoway requests) in order to prevent the client from waiting indefinitely for a response. However, the JVM may prevent Ice from successfully sending this response, which is another reason your clients should use invocation timeouts as a defensive strategy. Finally, in nearly all occurrences of an error, the server is unlikely to continue operating correctly even if Ice is able to complete the client's request. For example, the JVM terminates the thread that raised an uncaught error. In the case

Copyright 2016, ZeroC, Inc.

of an uncaught error raised by a servant, the thread being terminated is normally from the Ice server-side thread pool. If all of the threads in this pool eventually terminate due to uncaught errors, the server can no longer respond to new client requests.

See Also

- Run-Time Exceptions
- Java Mapping for Exceptions
- Server-Side Java Mapping for Interfaces
- Parameter Passing in Java
- Tie Classes in Java

Copyright 2016, ZeroC, Inc.

Tie Classes in Java

The mapping to skeleton classes requires the servant class to inherit from its skeleton class. Occasionally, this creates a problem: some class libraries require you to inherit from a base class in order to access functionality provided by the library; because Java does not support multiple implementation inheritance, this means that you cannot use such a class library to implement your servants because your servants cannot inherit from both the library class and the skeleton class simultaneously.

To allow you to still use such class libraries, Ice provides a way to write servants that replaces inheritance with delegation. This approach is supported by *tie classes*. The idea is that, instead of inheriting from the skeleton class, you simply create a class (known as an *implementation class* or *delegate class*) that contains methods corresponding to the operations of an interface. You use the `--tie` option with the `slice2java` compiler to create a tie class. For example, the `--tie` option causes the compiler to create exactly the same code for the `Node` interface as we saw previously, but to also emit an additional tie class. For an interface `<interface-name>`, the generated tie class has the name `_<interface-name>Tie`:

Java

Copyright 2016, ZeroC, Inc.

```
package Filesystem;

public class _NodeTie extends _NodeDisp implements Ice.TieBase {

    public _NodeTie() {}

    public
    _NodeTie(_NodeOperations delegate)
    {
        _ice_delegate = delegate;
    }

    public java.lang.Object
    ice_delegate()
    {
        return _ice_delegate;
    }

    public void
    ice_delegate(java.lang.Object delegate)
    {
        _ice_delegate = (_NodeOperations)delegate;
    }

    public boolean
    equals(java.lang.Object rhs)
    {
        if (this == rhs)
        {
            return true;
        }
        if (!(rhs instanceof _NodeTie))
        {
            return false;
        }

        return _ice_delegate.equals(((_NodeTie)rhs)._ice_delegate);
    }

    public int
    hashCode()
    {
        return _ice_delegate.hashCode();
    }

    public String
    name(Ice.Current current)
    {
        return _ice_delegate.name(current);
    }

    private _NodeOperations _ice_delegate;
}
```


Copyright 2016, ZeroC, Inc.

This looks a lot worse than it is: in essence, the generated tie class is simply a servant class (it extends `_NodeDisp`) that delegates to your implementation class each invocation of a method corresponding to a Slice operation:

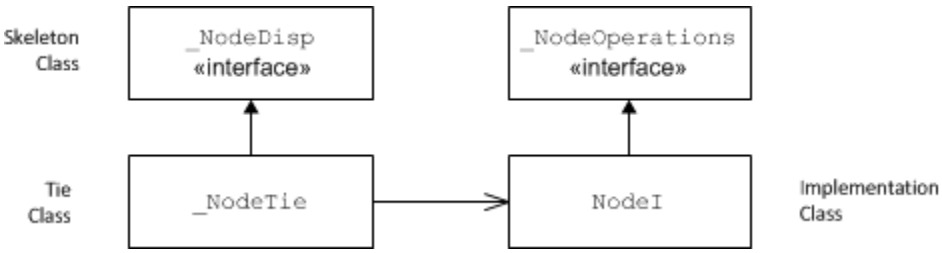

A skeleton class, tie class, and implementation class.

The generated tie class also implements the `Ice.TieBase` interface, which defines methods for obtaining and changing the delegate object:

Java

```
package Ice;

public interface TieBase {
    java.lang.Object ice_delegate();
    void ice_delegate(java.lang.Object delegate);
}
```

The delegate has type `java.lang.Object` in these methods in order to allow a tie object's delegate to be manipulated without knowing its actual type. However, the `ice_delegate` modifier raises `ClassCastException` if the given delegate object is not of the correct type.

Given this machinery, we can create an implementation class for our `Node` interface as follows:

Java

```
package Filesystem;

public final class NodeI implements _NodeOperations {

    public NodeI(String name)
    {
        _name = name;
    }

    public String name(Ice.Current current)
    {
        return _name;
    }

    private String _name;
}
```

Note that this class is identical to our previous implementation, except that it implements the `_NodeOperations` interface and does not extend `_NodeDisp` (which means that you are now free to extend any other class to support your implementation).

To create a servant, you instantiate your implementation class and the tie class, passing a reference to the implementation instance to the tie

Copyright 2016, ZeroC, Inc.

constructor:

Java

```
NodeI fred = new NodeI("Fred");         // Create implementation
_NodeTie servant = new _NodeTie(fred);  // Create tie
```

Alternatively, you can also default-construct the tie class and later set its delegate instance by calling `ice_delegate`:

Java

```
_NodeTie servant = new _NodeTie();      // Create tie
// ...
NodeI fred = new NodeI("Fred");         // Create implementation
// ...
servant.ice_delegate(fred);             // Set delegate
```

When using tie classes, it is important to remember that the tie instance is the servant, not your delegate. Furthermore, you must not use a tie instance to incarnate an Ice object until the tie has a delegate. Once you have set the delegate, you must not change it for the lifetime of the tie; otherwise, undefined behavior results.

You should use the tie approach only when necessary, that is, if you need to extend some base class in order to implement your servants: using the tie approach is more costly in terms of memory because each Ice object is incarnated by two Java objects (the tie and the delegate) instead of just one. In addition, call dispatch for ties is marginally slower than for ordinary servants because the tie forwards each operation to the delegate, that is, each operation invocation requires two function calls instead of one.

Also note that, unless you arrange for it, there is no way to get from the delegate back to the tie. If you need to navigate back to the tie from the delegate, you can store a reference to the tie in a member of the delegate. (The reference can, for example, be initialized by the constructor of the delegate.)

See Also

- Server-Side Java Mapping for Interfaces
- Parameter Passing in Java
- Raising Exceptions in Java
- Object Incarnation in Java

Copyright 2016, ZeroC, Inc.

Object Incarnation in Java

Having created a servant class such as the rudimentary `NodeI` class, you can instantiate the class to create a concrete servant that can receive invocations from a client. However, merely instantiating a servant class is insufficient to incarnate an object. Specifically, to provide an implementation of an Ice object, you must take the following steps:

1. Instantiate a servant class.
2. Create an identity for the Ice object incarnated by the servant.
3. Inform the Ice run time of the existence of the servant.
4. Pass a proxy for the object to a client so the client can reach it.

On this page:

- Instantiating a Java Servant
- Creating an Identity in Java
- Activating a Java Servant
- UUIDs as Identities in Java
- Creating Proxies in Java
  - Proxies and Servant Activation in Java
  - Direct Proxy Creation in Java

Instantiating a Java Servant

Instantiating a servant means to allocate an instance:

Java

```
Node servant = new NodeI("Fred");
```

This code creates a new `NodeI` instance and assigns its address to a reference of type `Node`. This works because `NodeI` is derived from `Node`, so a `Node` reference can refer to an instance of type `NodeI`. However, if we want to invoke a member function of the `NodeI` class at this point, we must use a `NodeI` reference:

Java

```
NodeI servant = new NodeI("Fred");
```

Whether you use a `Node` or a `NodeI` reference depends purely on whether you want to invoke a member function of the `NodeI` class: if not, a `Node` reference works just as well as a `NodeI` reference.

Creating an Identity in Java

Each Ice object requires an identity. That identity must be unique for all servants using the same object adapter.

The Ice object model assumes that all objects (regardless of their adapter) have a globally unique identity.

An Ice object identity is a structure with the following Slice definition:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    struct Identity {
        string name;
        string category;
    };
    // ...
};
```

The full identity of an object is the combination of both the `name` and `category` fields of the `Identity` structure. For now, we will leave the `category` field as the empty string and simply use the `name` field. (The `category` field is most often used in conjunction with servant locators.)

To create an identity, we simply assign a key that identifies the servant to the `name` field of the `Identity` structure:

Java

```
Ice.Identity id = new Ice.Identity();
id.name = "Fred"; // Not unique, but good enough for now
```

Activating a Java Servant

Merely creating a servant instance does nothing: the Ice run time becomes aware of the existence of a servant only once you explicitly tell the object adapter about the servant. To activate a servant, you invoke the `add` operation on the object adapter. Assuming that we have access to the object adapter in the `_adapter` variable, we can write:

Java

```
_adapter.add(servant, id);
```

Note the two arguments to `add`: the servant and the object identity. Calling `add` on the object adapter adds the servant and the servant's identity to the adapter's servant map and links the proxy for an Ice object to the correct servant instance in the server's memory as follows:

1. The proxy for an Ice object, apart from addressing information, contains the identity of the Ice object. When a client invokes an operation, the object identity is sent with the request to the server.
2. The object adapter receives the request, retrieves the identity, and uses the identity as an index into the servant map.
3. If a servant with that identity is active, the object adapter retrieves the servant from the servant map and dispatches the incoming request into the correct member function on the servant.

Assuming that the object adapter is in the active state, client requests are dispatched to the servant as soon as you call `add`.

UUIDs as Identities in Java

The Ice object model assumes that object identities are globally unique. One way of ensuring that uniqueness is to use UUIDs (Universally Unique Identifiers) as identities. Java provides a helper function that we can use to create such identities:

Copyright 2016, ZeroC, Inc.

Java

```
public class Example {
    public static void
    main(String[] args)
    {
        System.out.println(java.util.UUID.randomUUID().toString());
    }
}
```

When executed, this program prints a unique string such as `5029a22c-e333-4f87-86b1-cd5e0fcce509`. Each call to `randomUUID` creates a string that differs from all previous ones.

You can use a UUID such as this to create object identities. For convenience, the object adapter has an operation `addWithUUID` that generates a UUID and adds a servant to the servant map in a single step. Using this operation, we can create an identity and register a servant with that identity in a single step as follows:

Java

```
_adapter.addWithUUID(new NodeI("Fred"));
```

Creating Proxies in Java

Once we have activated a servant for an Ice object, the server can process incoming client requests for that object. However, clients can only access the object once they hold a proxy for the object. If a client knows the server's address details and the object identity, it can create a proxy from a string, as we saw in our first example in Hello World Application. However, creation of proxies by the client in this manner is usually only done to allow the client access to initial objects for bootstrapping. Once the client has an initial proxy, it typically obtains further proxies by invoking operations.

The object adapter contains all the details that make up the information in a proxy: the addressing and protocol information, and the object identity. The Ice run time offers a number of ways to create proxies. Once created, you can pass a proxy to the client as the return value or as an out-parameter of an operation invocation.

Proxies and Servant Activation in Java

The `add` and `addWithUUID` servant activation operations on the object adapter return a proxy for the corresponding Ice object. This means we can write:

Java

```
NodePrx proxy = NodePrxHelper.uncheckedCast(_adapter.addWithUUID(new NodeI("Fred"))
);
```

Here, `addWithUUID` both activates the servant and returns a proxy for the Ice object incarnated by that servant in a single step.

Note that we need to use an `uncheckedCast` here because `addWithUUID` returns a proxy of type `Ice.ObjectPrx`.

Direct Proxy Creation in Java

The object adapter offers an operation to create a proxy for a given identity:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local interface ObjectAdapter {
        Object* createProxy(Identity id);
        // ...
    };
};
```

Note that `createProxy` creates a proxy for a given identity whether a servant is activated with that identity or not. In other words, proxies have a life cycle that is quite independent from the life cycle of servants:

Java

```
Ice.Identity id = new Ice.Identity();
id.name = java.util.UUID.randomUUID().toString();
Ice.ObjectPrx o = _adapter.createProxy(id);
```

This creates a proxy for an Ice object with the identity returned by `randomUUID`. Obviously, no servant yet exists for that object so, if we return the proxy to a client and the client invokes an operation on the proxy, the client will receive an `ObjectNotExistException`. (We examine these life cycle issues in more detail in Object Life Cycle.)

See Also

- Hello World Application
- Server-Side Java Mapping for Interfaces
- Object Adapter States
- Servant Locators
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

Asynchronous Method Dispatch (AMD) in Java

The number of simultaneous synchronous requests a server is capable of supporting is determined by the number of threads in the server's t hread pool. If all of the threads are busy dispatching long-running operations, then no threads are available to process new requests and therefore clients may experience an unacceptable lack of responsiveness.

Asynchronous Method Dispatch (AMD), the server-side equivalent of AMI, addresses this scalability issue. Using AMD, a server can receive a request but then suspend its processing in order to release the dispatch thread as soon as possible. When processing resumes and the results are available, the server sends a response explicitly using a callback object provided by the Ice run time.

AMD is transparent to the client, that is, there is no way for a client to distinguish a request that, in the server, is processed synchronously from a request that is processed asynchronously.

In practical terms, an AMD operation typically queues the request data (i.e., the callback object and operation arguments) for later processing by an application thread (or thread pool). In this way, the server minimizes the use of dispatch threads and becomes capable of efficiently supporting thousands of simultaneous clients.

An alternate use case for AMD is an operation that requires further processing after completing the client's request. In order to minimize the client's delay, the operation returns the results while still in the dispatch thread, and then continues using the dispatch thread for additional work.

On this page:

- Enabling AMD with Metadata in Java
- AMD Mapping in Java
  - Callback interface for AMD
  - Dispatch method for AMD
- AMD Exceptions in Java
- AMD Example in Java

Enabling AMD with Metadata in Java

To enable asynchronous dispatch, you must add an `["amd"]` metadata directive to your Slice definitions. The directive applies at the interface and the operation level. If you specify `["amd"]` at the interface level, all operations in that interface use asynchronous dispatch; if you specify `["amd"]` for an individual operation, only that operation uses asynchronous dispatch. In either case, the metadata directive *repl aces* synchronous dispatch, that is, a particular operation implementation must use synchronous or asynchronous dispatch and cannot use both.

Consider the following Slice definitions:

Slice

```
["amd"] interface I {
    bool isValid();
    float computeRate();
};

interface J {
    ["amd"] void startProcess();
    int endProcess();
};
```

In this example, both operations of interface `I` use asynchronous dispatch, whereas, for interface `J`, `startProcess` uses asynchronous dispatch and `endProcess` uses synchronous dispatch.

Specifying metadata at the operation level (rather than at the interface or class level) minimizes the amount of generated code and, more importantly, minimizes complexity: although the asynchronous model is more flexible, it is also more complicated to use. It is therefore in your best interest to limit the use of the asynchronous model to those operations that need it, while using the simpler synchronous model for the rest.

AMD Mapping in Java

The Java mapping emits the following code for each AMD operation:

Copyright 2016, ZeroC, Inc.

1. Callback interface
2. Dispatch method

Callback interface for AMD

A callback interface is used by the implementation to notify the Ice run time about the completion of an operation. The name of this interface is formed using the pattern `AMD_class_op`. For example, an operation named `foo` defined in interface `I` results in an interface named `AMD_I_foo`. The interface is generated in the same scope as the interface or class containing the operation. Two methods are provided:

Java

```
public void ice_response(<params>);
```

The `ice_response` method allows the server to report the successful completion of the operation. If the operation has a non-`void` return type, the first parameter to `ice_response` is the return value. Parameters corresponding to the operation's out parameters follow the return value, in the order of declaration.

Java

```
public void ice_exception(java.lang.Exception ex);
```

The `ice_exception` method allows the server to raise an exception. With respect to exceptions, there is less compile-time type safety in an AMD implementation because there is no `throws` clause on the dispatch method and any exception type could conceivably be passed to `ice_exception`. However, the Ice run time validates the exception value using the same semantics as for synchronous dispatch.

Neither `ice_response` nor `ice_exception` throw any exceptions to the caller.

Dispatch method for AMD

The dispatch method, whose name has the suffix `_async`, has a `void` return type. The first parameter is a reference to an instance of the callback interface described above. The remaining parameters comprise the `in` parameters of the operation, in the order of declaration.

For example, suppose we have defined the following operation:

Slice

```
interface I {
  ["amd"] int foo(short s, out long l);
};
```

The callback interface generated for operation `foo` is shown below:

Copyright 2016, ZeroC, Inc.

Java

```
public interface AMD_I_foo {
    void ice_response(int __ret, long l);
    void ice_exception(java.lang.Exception ex);
}
```

The dispatch method for asynchronous invocation of operation `foo` is generated as follows:

Java

```
void foo_async(AMD_I_foo __cb, short s);
```

AMD Exceptions in Java

There are two processing contexts in which the logical implementation of an AMD operation may need to report an exception: the dispatch thread (the thread that receives the invocation), and the response thread (the thread that sends the response).

These are not necessarily two different threads: it is legal to send the response from the dispatch thread.

Although we recommend that the callback object be used to report all exceptions to the client, it is legal for the implementation to raise an exception instead, but only from the dispatch thread.

As you would expect, an exception raised from a response thread cannot be caught by the Ice run time; the application's run time environment determines how such an exception is handled. Therefore, a response thread must ensure that it traps all exceptions and sends the appropriate response using the callback object. Otherwise, if a response thread is terminated by an uncaught exception, the request may never be completed and the client might wait indefinitely for a response.

Whether raised in a dispatch thread or reported via the callback object, user exceptions are validated and local exceptions may undergo translation.

AMD Example in Java

To demonstrate the use of AMD in Ice, let us define the Slice interface for a simple computational engine:

Slice

```
module Demo {
    sequence<float> Row;
    sequence<Row> Grid;

    exception RangeError {};

    interface Model {
        ["amd"] Grid interpolate(Grid data, float factor)
            throws RangeError;
    };
};
```


Copyright 2016, ZeroC, Inc.

Given a two-dimensional grid of floating point values and a factor, the `interpolate` operation returns a new grid of the same size with the values interpolated in some interesting (but unspecified) way.

Our servant class derives from `Demo._ModelDisp` and supplies a definition for the `interpolate_async` method that creates a `Job` to hold the callback object and arguments, and adds the `Job` to a queue. The method is synchronized to guard access to the queue:

Java

```
public final class ModelI extends Demo._ModelDisp {
    synchronized public void interpolate_async(
        Demo.AMD_Model_interpolate cb,
        float[][] data,
        float factor,
        Ice.Current current)
            throws RangeError
    {
        _jobs.add(new Job(cb, data, factor));
    }

    java.util.LinkedList<Job> _jobs = new java.util.LinkedList<Job>();
}
```

After queuing the information, the operation returns control to the Ice run time, making the dispatch thread available to process another request. An application thread removes the next `Job` from the queue and invokes `execute`, which uses `interpolateGrid` (not shown) to perform the computational work:

Copyright 2016, ZeroC, Inc.

Java

```
class Job {
    Job(Demo.AMD_Model_interpolate cb,
        float[][] grid,
        float factor)
    {
        _cb = cb;
        _grid = grid;
        _factor = factor;
    }

    void execute()
    {
        if (!interpolateGrid()) {
            _cb.ice_exception(new Demo.RangeError());
            return;
        }
        _cb.ice_response(_grid);
    }

    private boolean interpolateGrid() {
        // ...
    }

    private Demo.AMD_Model_interpolate _cb;
    private float[][] _grid;
    private float _factor;
}
```

If `interpolateGrid` returns `false`, then `ice_exception` is invoked to indicate that a range error has occurred. The `return` statement following the call to `ice_exception` is necessary because `ice_exception` does not throw an exception; it only marshals the exception argument and sends it to the client.

If interpolation was successful, `ice_response` is called to send the modified grid back to the client.

See Also

- User Exceptions
- Run-Time Exceptions
- Asynchronous Method Invocation (AMI) in Java
- The Ice Threading Model

Copyright 2016, ZeroC, Inc.

Example of a File System Server in Java

This page presents the source code for a Java server that implements our file system and communicates with the client we wrote earlier. The code is fully functional, apart from the required interlocking for threads.

Note that the server is remarkably free of code that relates to distribution: most of the server code is simply application logic that would be present just the same for a non-distributed version. Again, this is one of the major advantages of Ice: distribution concerns are kept away from application code so that you can concentrate on developing application logic instead of networking infrastructure.

The server code shown here is not quite correct as it stands: if two clients access the same file in parallel, each via a different thread, one thread may read the `_lines` data member while another thread updates it. Obviously, if that happens, we may write or return garbage or, worse, crash the server. However, it is trivial to make the `read` and `write` operations thread-safe. We discuss thread safety in The Ice Threading Model.

On this page:

- Implementing a File System Server in Java
- Server Main Program in Java
- FileI Servant Class in Java
- DirectoryI Servant Class in Java
  - DirectoryI Data Members
  - DirectoryI Constructor
  - DirectoryI Methods

Implementing a File System Server in Java

We have now seen enough of the server-side Java mapping to implement a server for our file system. (You may find it useful to review these Slice definitions before studying the source code.)

Our server is composed of three source files:

- `Server.java`
  This file contains the server main program.
- `Filesystem/DirectoryI.java`
  This file contains the implementation for the `Directory` servants.
- `Filesystem/FileI.java`
  This file contains the implementation for the `File` servants.

Server Main Program in Java

Our server main program, in the file `Server.java`, uses the `Ice.Application` class. The `run` method installs a shutdown hook, creates an object adapter, instantiates a few servants for the directories and files in the file system, and then activates the adapter. This leads to a main program as follows:

Java

```
import Filesystem.*;

public class Server extends Ice.Application {
    public int
    run(String[] args)
    {
        //
        // Terminate cleanly on receipt of a signal
        //
        shutdownOnInterrupt();

        // Create an object adapter (stored in the _adapter
        // static members)
```


Copyright 2016, ZeroC, Inc.

```
        //
        Ice.ObjectAdapter adapter
= communicator().createObjectAdapterWithEndpoints(
            "SimpleFilesystem", "default -p 10000");
        DirectoryI._adapter = adapter;
        FileI._adapter = adapter;

        // Create the root directory (with name "/" and no parent)
        //
        DirectoryI root = new DirectoryI("/", null);

        // Create a file "README" in the root directory
        //
        File file = new FileI("README", root);
        String[] text;
        text = new String[] {
            "This file system contains a collection of poetry."
        };
        try {
            file.write(text, null);
        } catch (GenericError e) {
            System.err.println(e.reason);
        }

        // Create a directory "Coleridge" in the root directory
        //
        DirectoryI coleridge = new DirectoryI("Coleridge", root);

        // Create a file "Kubla_Khan" in the Coleridge directory
        //
        file = new FileI("Kubla_Khan", coleridge);
        text = new String[]{ "In Xanadu did Kubla Khan",
                             "A stately pleasure-dome decree:",
                             "Where Alph, the sacred river, ran",
                             "Through caverns measureless to man",
                             "Down to a sunless sea." };
        try {
            file.write(text, null);
        } catch (GenericError e) {
            System.err.println(e.reason);
        }

        // All objects are created, allow client requests now
        //
        adapter.activate();

        // Wait until we are done
        //
        communicator().waitForShutdown();

        return 0;
    }

    public static void
    main(String[] args)
    {
```


Copyright 2016, ZeroC, Inc.

```
        Server app = new Server();
        System.exit(app.main("Server", args));
```


Copyright 2016, ZeroC, Inc.

```
    }
}
```

The code imports the contents of the `Filesystem` package. This avoids having to continuously use fully-qualified identifiers with a `Filesystem.` prefix.

The next part of the source code is the definition of the `Server` class, which derives from `Ice.Application` and contains the main application logic in its `run` method. Much of this code is boiler plate that we saw previously: we create an object adapter, and, towards the end, activate the object adapter and call `waitForShutdown`.

The interesting part of the code follows the adapter creation: here, the server instantiates a few nodes for our file system to create the structure shown below:

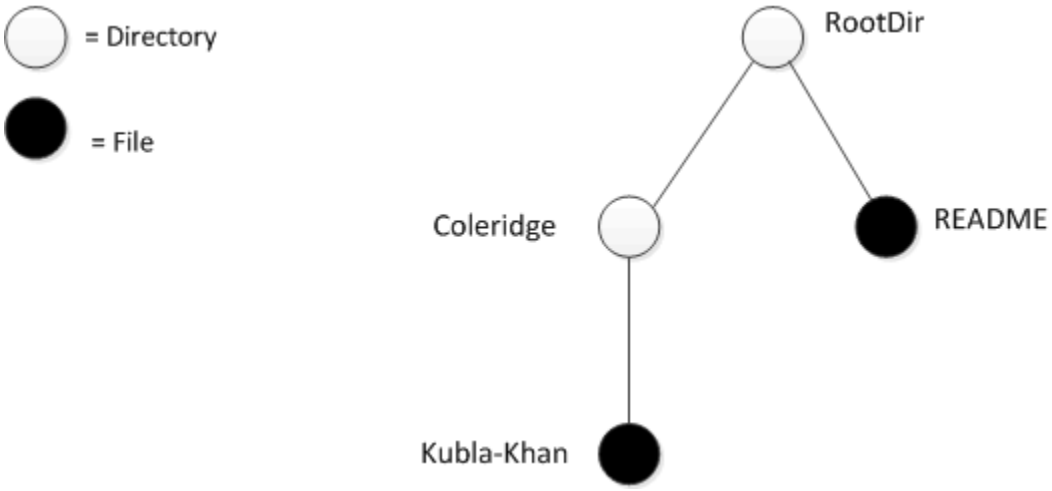

A small file system.

As we will see shortly, the servants for our directories and files are of type `DirectoryI` and `FileI`, respectively. The constructor for either type of servant accepts two parameters, the name of the directory or file to be created and a reference to the servant for the parent directory. (For the root directory, which has no parent, we pass a null parent.) Thus, the statement

Java

```
DirectoryI root = new DirectoryI("/", null);
```

creates the root directory, with the name `"/"` and no parent directory.

Here is the code that establishes the structure in the above illustration:

Copyright 2016, ZeroC, Inc.

Java

```
// Create the root directory (with name "/" and no parent)
//
DirectoryI root = new DirectoryI("/", null);

// Create a file "README" in the root directory
//
File file = new FileI("README", root);
String[] text;
text = new String[] {
    "This file system contains a collection of poetry."
};
try {
    file.write(text, null);
} catch (GenericError e) {
    System.err.println(e.reason);
}

// Create a directory "Coleridge" in the root directory
//
DirectoryI coleridge = new DirectoryI("Coleridge", root);

// Create a file "Kubla_Khan" in the Coleridge directory
//
file = new FileI("Kubla_Khan", coleridge);
text = new String[]{ "In Xanadu did Kubla Khan",
                     "A stately pleasure-dome decree:",
                     "Where Alph, the sacred river, ran",
                     "Through caverns measureless to man",
                     "Down to a sunless sea." };
try {
    file.write(text, null);
} catch (GenericError e) {
    System.err.println(e.reason);
}
```

We first create the root directory and a file `README` within the root directory. (Note that we pass a reference to the root directory as the parent when we create the new node of type `FileI`.)

The next step is to fill the file with text:

Copyright 2016, ZeroC, Inc.

Java

```
        String[] text;
        text = new String[] {
            "This file system contains a collection of poetry."
        };
        try {
            file.write(text, null);
        } catch (GenericError e) {
            System.err.println(e.reason);
        }
```

Recall that Slice sequences by default map to Java arrays. The Slice type `Lines` is simply an array of strings; we add a line of text to our `README` file by initializing the `text` array to contain one element.

Finally, we call the Slice `write` operation on our `FileI` servant by writing:

Java

```
            file.write(text, null);
```

This statement is interesting: the server code invokes an operation on one of its own servants. Because the call happens via a reference to the servant (of type `FileI`) and *not* via a proxy (of type `FilePrx`), the Ice run time does not know that this call is even taking place — such a direct call into a servant is not mediated by the Ice run time in any way and is dispatched as an ordinary Java function call.

In similar fashion, the remainder of the code creates a subdirectory called `Coleridge` and, within that directory, a file called `Kubla_Khan` to complete the structure in the illustration listed above.

`FileI` Servant Class in Java

Our `FileI` servant class has the following basic structure:

Java

```
public class FileI extends _FileDisp
{
    // Constructor and operations here...

    public static Ice.ObjectAdapter _adapter;
    private String _name;
    private DirectoryI _parent;
    private String[] _lines;
}
```

The class has a number of data members:

- `_adapter`
  This static member stores a reference to the single object adapter we use in our server.

Copyright 2016, ZeroC, Inc.

- `_name`
  This member stores the name of the file incarnated by the servant.
- `_parent`
  This member stores the reference to the servant for the file's parent directory.
- `_lines`
  This member holds the contents of the file.

The `_name` and `_parent` data members are initialized by the constructor:

Java

```
public
FileI(String name, DirectoryI parent)
{
    _name = name;
    _parent = parent;

    assert(_parent != null);

    // Create an identity
    //
    Ice.Identity myID = new Ice.Identity();
    myID.name = java.util.UUID.randomUUID().toString();

    // Add the identity to the object adapter
    //
    _adapter.add(this, myID);

    // Create a proxy for the new node and
    // add it as a child to the parent
    //
    NodePrx thisNode = NodePrxHelper.uncheckedCast(_adapter.createProxy(myID));
    _parent.addChild(thisNode);
}
```

After initializing the `_name` and `_parent` members, the code verifies that the reference to the parent is not null because every file must have a parent directory. The constructor then generates an identity for the file by calling `java.util.UUID.randomUUID` and adds itself to the servant map by calling `ObjectAdapter.add`. Finally, the constructor creates a proxy for this file and calls the `addChild` method on its parent directory. `addChild` is a helper function that a child directory or file calls to add itself to the list of descendant nodes of its parent directory. We will see the implementation of this function in `DirectoryI` Methods.

The remaining methods of the `FileI` class implement the Slice operations we defined in the `Node` and `File` Slice interfaces:

Copyright 2016, ZeroC, Inc.

Java

```
// Slice Node::name() operation

public String
name(Ice.Current current)
{
    return _name;
}

// Slice File::read() operation

public String[]
read(Ice.Current current)
{
    return _lines;
}

// Slice File::write() operation

public void
write(String[] text, Ice.Current current)
    throws GenericError
{
    _lines = text;
}
```

The `name` method is inherited from the generated `Node` interface (which is a base interface of the `_FileDisp` class from which `FileI` is derived). It returns the value of the `_name` member.

The `read` and `write` methods are inherited from the generated `File` interface (which is a base interface of the `_FileDisp` class from which `FileI` is derived) and return and set the `_lines` member.

`DirectoryI` Servant Class in Java

The `DirectoryI` class has the following basic structure:

Copyright 2016, ZeroC, Inc.

Java

```
package Filesystem;

public final class DirectoryI extends _DirectoryDisp
{
    // Constructor and operations here...

    public static Ice.ObjectAdapter _adapter;
    private String _name;
    private DirectoryI _parent;
    private java.util.ArrayList<NodePrx> _contents
= new java.util.ArrayList<NodePrx>();
}
```

`DirectoryI` Data Members

As for the `FileI` class, we have data members to store the object adapter, the name, and the parent directory. (For the root directory, the `_parent` member holds a null reference.) In addition, we have a `_contents` data member that stores the list of child directories. These data members are initialized by the constructor:

Java

```
public
DirectoryI(String name, DirectoryI parent)
{
    _name = name;
    _parent = parent;

    // Create an identity. The parent has the
    // fixed identity "RootDir"
    //
    Ice.Identity myID = new Ice.Identity();
    myID.name = _parent != null ? java.util.UUID.randomUUID().toString() : "RootDir
";

    // Add the identity to the object adapter
    //
    _adapter.add(this, myID);

    // Create a proxy for the new node and add it as a
    // child to the parent
    //
    NodePrx thisNode = NodePrxHelper.uncheckedCast(_adapter.createProxy(myID));
    if (_parent != null)
        _parent.addChild(thisNode);
}
```

`DirectoryI` Constructor

Copyright 2016, ZeroC, Inc.

The constructor creates an identity for the new directory by calling `java.util.UUID.randomUUID`. (For the root directory, we use the fixed identity `"RootDir"`.) The servant adds itself to the servant map by calling `ObjectAdapter.add` and then creates a reference to itself and passes it to the `addChild` helper function.

`DirectoryI` Methods

`addChild` adds the passed reference to the `_contents` list:

Java

```
void
addChild(NodePrx child)
{
    _contents.add(child);
}
```

The remainder of the operations, `name` and `list`, are equally trivial:

Java

```
public String
name(Ice.Current current)
{
    return _name;
}

// Slice Directory::list() operation

public NodePrx[]
list(Ice.Current current)
{
    NodePrx[] result = new NodePrx[_contents.size()];
    _contents.toArray(result);
    return result;
}
```

Note that the `_contents` member is of type `java.util.ArrayList<NodePrx>`, which is convenient for the implementation of the `addChild` method. However, this requires us to convert the list into a Java array in order to return it from the `list` operation.

See Also

- Slice for a Simple File System
- Example of a File System Client in Java
- The Server-Side main Method in Java
- Java Mapping for Sequences
- The Ice Threading Model

Copyright 2016, ZeroC, Inc.

The Java Utility Library

Ice for Java includes a number of utility APIs in the `IceUtil` package and the `Ice.Util` class. This section summarizes the contents of these APIs for your reference.

On this page:

- The IceUtil Package in Java
  - Cache and Store Classes
- The Ice.Util Class in Java
  - Communicator Initialization Methods
  - Identity Conversion
  - Per-Process Logger Methods
  - Property Creation Methods
  - Proxy Comparison Methods
  - Stream Creation
  - Version Information

The `IceUtil` Package in Java

`Cache` and `Store` Classes

The `Cache` class allows you to efficiently maintain a cache that is backed by secondary storage, such as a Berkeley DB database, without holding a lock on the entire cache while values are being loaded from the database. If you want to create evictors for servants that store their state in a database, the `Cache` class can simplify your evictor implementation considerably.

You may also want to examine the implementation of the Freeze background save evictor in the source distribution; it uses `IceUtil.Cache` for its implementation.

The `Cache` class has the following interface:

Java

```
package IceUtil;

public class Cache {
    public Cache(Store store);

    public Object pin(Object key);
    public Object pin(Object key, Object o);
    public Object unpin(Object key);

    public Object putIfAbsent(Object key, Object newObj);
    public Object getIfPinned(Object key);

    public void clear();
    public int size();
}
```

Internally, a `Cache` maintains a map of name-value pairs. The implementation of `Cache` takes care of maintaining the map; in particular, it ensures that concurrent lookups by callers are possible without blocking even if some of the callers are currently loading values from the backing store. In turn, this is useful for evictor implementations, such as the Freeze background save evictor. The `Cache` class does not limit the number of entries in the cache — it is the job of the evictor implementation to limit the map size by calling `unpin` on elements of the map that it wants to evict.

The `Cache` class works in conjunction with a `Store` interface for which you must provide an implementation. The `Store` interface is trivial:

Copyright 2016, ZeroC, Inc.

Java

```
package IceUtil;

public interface Store {
    Object load(Object key);
}
```

You must implement the `load` method in a class that you derive from `Store`. The `Cache` implementation calls `load` when it needs to retrieve the value for the passed key from the backing store. If `load` cannot locate a record for the given key because no such record exists, it must return null. If `load` fails for some other reason, it can throw an exception derived from `java.lang.RuntimeException`, which is propagated back to the application code.

The public member functions of `Cache` behave as follows:
`Cache(Store s)`

The constructor initializes the cache with your implementation of the `Store` interface.
`Object pin(Object key, Object val)`

To add a key-value pair to the cache, your evictor can call `pin`. The return value is null if the key and value were added; otherwise, if the map already contains an entry with the given key, the entry is unchanged and `pin` returns the original value for that key.

This version of `pin` does *not* call `load` to retrieve the entry from backing store if it is not yet in the cache. This is useful when you add a newly-created object to the cache.
`Object pin(Object key)`

This version of `pin` returns the value stored in the cache for the given key if the cache already contains an entry for that key. If no entry with the given key is in the cache, `pin` calls `load` to retrieve the corresponding value (if any) from the backing store. `pin` returns the value returned by `load`, that is, the value if `load` could retrieve it, null if `load` could not retrieve it, or any exception thrown by `load`.
`Object unpin(Object key)`

`unpin` removes the entry for the given key from the cache. If the cache contained an entry for the key, the return value is the value for that key; otherwise, the return value is null.
`Object putIfAbsent(Object key, Object val)`

This function adds a key-value pair to the cache. If the cache already contains an entry for the given key, `putIfAbsent` returns the original value for that key. If no entry with the given key is in the cache, `putIfAbsent` calls `load` to retrieve the corresponding entry (if any) from the backing store and returns the value returned by `load`.

If the cache does not contain an entry for the given key and `load` does not retrieve a value for the key, the method adds the new entry and returns null.
`Object getIfPinned(Object key)`

This function returns the value stored for the given key. If an entry for the given key is in the map, the function returns the corresponding value; otherwise, the function returns null. `getIfPinned` does not call `load`.
`void clear()`

This function removes all entries in the map.
`int size()`

This function returns the number of entries in the map.

The `Ice.Util` Class in Java

Communicator Initialization Methods

`Ice.Util` provides a number of overloaded `initialize` methods that create a communicator.

Identity Conversion

Copyright 2016, ZeroC, Inc.

`Ice.Util` contains two methods for converting object identities of type `Ice.Identity` to and from strings.

Per-Process Logger Methods

`Ice.Util` provides methods for getting and setting the per-process logger.

Property Creation Methods

`Ice.Util` provides a number of overloaded `createProperties` methods that create property sets.

Proxy Comparison Methods

Two methods, `proxyIdentityCompare` and `proxyIdentityAndFacetCompare`, allow you to compare object identities that are stored in proxies (either ignoring the facet or taking the facet into account).

Stream Creation

The methods `createInputStream`, `wrapInputStream` and `createOutputStream` create streams for use with dynamic invocation.

Version Information

The `stringVersion` and `intVersion` methods return the version of the Ice run time:

Java

```
public static String stringVersion();
public static int intVersion();
```

The `stringVersion` method returns the Ice version in the form `<major>.<minor>.<patch>`, for example, `3.4.2`. For beta releases, the version is `<major>.<minor>`b, for example, `3.4b`.

The `intVersion` method returns the Ice version in the form *AABBCC*, where *AA* is the major version number, *BB* is the minor version number, and *CC* is patch level, for example, `30402` for version 3.4.2. For beta releases, the patch level is set to 51 so, for example, for version 3.4b, the value is `30451`.

See Also

- Background Save Evictor
- Java Mapping for Interfaces
- Command-Line Parsing and Initialization
- Setting Properties
- Object Identity
- Java Streaming Interfaces

Copyright 2016, ZeroC, Inc.

Custom Class Loaders

Certain features of the Ice for Java run-time necessitate dynamic class loading. Applications with special requirements can supply a custom class loader for Ice to use in the following situations:

- Unmarshaling user exceptions and instances of concrete Slice classes
- Loading Ice plug-ins
- Loading IceSSL certificate verifiers and password callbacks

If an application does not supply a class loader (or if the application-supplied class loader fails to locate a class), the Ice run time attempts to load the class using class loaders in the following order:

- current thread's class loader
- default class loader (that is, by calling `Class.forName`)
- system class loader

Note that an application must install object factories for any abstract Slice classes it might receive, regardless of whether the application also installs a custom class loader.

To install a custom class loader, set the `classLoader` member of `Ice.InitializationData` prior to creating a communicator:

Java

```
Ice.InitializationData initData = new Ice.InitializationData();
initData.classLoader = new MyClassLoader();
Ice.Communicator communicator = Ice.Util.initialize(args, initData);
```

See Also

- Java Mapping for Exceptions
- Java Mapping for Classes
- Plug-in Facility
- IceSSL
- Communicator Initialization

Copyright 2016, ZeroC, Inc.

Java Interrupts

As of Ice 3.6, Java applications can safely interrupt a thread that calls a blocking Ice API.

> ⚠ You must enable the `Ice.ThreadInterruptSafe` property to use interrupts. Enabling interrupts causes Ice to disable message buffer caching, which may incur a slight performance penalty.

With interrupt support enabled, Ice guarantees that every call into the run time that could potentially block indefinitely is an *interruption point*. On entry, an interruption point raises the unchecked exception `Ice.OperationInterruptedException` immediately if the current thread's interrupted flag is true. If the application interrupts the run time after the interruption point has begun its processing, the interruption point will either raise `Ice.OperationInterruptedException` or return control to the application with the thread's interrupted flag set. Under no circumstances will Ice silently discard an interrupt.

For example, the following code shows how to interrupt a thread that's blocked while making a synchronous proxy invocation:

Java

```
Thread proxyThread = Thread.currentThread();
AccountPrx account = // get proxy...

try {
    String name = account.getName(); // Blocks calling thread
    ...
} catch (Ice.OperationInterruptedException ex) {
   // The invocation was interrupted.
}

// At this point either the invocation was interrupted with an
// OperationInterruptedException or proxyThread.isInterrupted() is true.

// In another thread either just prior to or during the getName() invocation.
proxyThread.interrupt();
```

Here is a similar example that uses the asynchronous proxy API instead:

Copyright 2016, ZeroC, Inc.

Java

```
Thread proxyThread = Thread.currentThread();
AccountPrx account = // get proxy...

Ice.AsyncResult r = account.begin_getName(); // Never blocks

try {
    // do more work
    String name = account.end_getName(r); // May block calling thread
    ...
} catch (Ice.OperationInterruptedException ex) {
    // The invocation was interrupted.
}

// At this point either the invocation was interrupted with an
// OperationInterruptedException or proxyThread.isInterrupted() is true.

// In another thread either just prior to or during the getName() invocation.
proxyThread.interrupt();
```

Synchronous proxy invocations are interruption points. For an asynchronous proxy invocation, the `begin_` method never blocks and therefore it is not an interruption point, however the `end_` method is an interruption point.

Additionally, all of the methods listed in Blocking API Calls are also interruption points. In the specific case of `Communicator.destroy()`, we recommend calling `destroy` in a loop as shown below:

Java

```
while(true) {
  try {
    communicator.destroy();
    break;
  } catch(Ice.OperationInterruptedException ex) {
    continue;
  }
}
```

Interrupting a call to `destroy` could leave the Ice run time in a partially-destroyed state; calling `destroy` again as we've done here ensures Ice can finish reclaiming the communicator's resources.

See Also

- Blocking API Calls
- Java Mapping

Copyright 2016, ZeroC, Inc.

JavaScript Mapping

Supported Features of the JavaScript Mapping

Ice for JavaScript does not support the entire set of Ice features that are found in the C++, Java, and C# language mappings, primarily due to platform and API limitations. Furthermore, there are some differences between using Ice for JavaScript in a browser and in Node.js.

The table below provides more details on the Ice for JavaScript feature set:

Feature	Browser	Node.js	Notes
Synchronous invocations			Blocking RPCs are not compatible with JavaScript's execution model.
Asynchronous invocations	✓	✓	Also see Asynchronous APIs in JavaScript
Synchronous dispatch	✓	✓	
Asynchronous dispatch	✓	✓	
Outgoing TCP connections		✓	
Incoming TCP connections			Applications must use bidirectional connections to receive callbacks.
Outgoing WS connections	✓		WebSocket connections require the server to use the IceWS transport.
Outgoing WSS connections	✓		Secure WebSocket (WSS) connections require the server to use the IceWS transport.
Outgoing SSL connections			Browser applications can use WSS.
Datagrams			
DNS queries	⚠	⚠	The lack of support for DNS queries affects fault tolerance.
File logging			
Property files			
Flow control API			"Sent" callbacks are not supported.
Thread pools			Threads are not supported in JavaScript.
`PerThread` implicit context			Threads are not supported in JavaScript.
Protocol compression			Ice uses the bzip2 algorithm, which is not natively supported in JavaScript.
Communicator plug-ins			
`Dispatcher` API			
`DispatchInterceptor` API			
Collocated invocations			
Streaming API			
Admin facility			

DNS limitations

In most language mappings, the Ice run time performs a DNS query to resolve an endpoint's host name into one or more IP addresses. Specifying a multi-homed host name in an endpoint provides the client with a simple form of fault tolerance because Ice has multiple addresses to use during connection establishment.

This form of fault tolerance is not available when using Ice for JavaScript because the DNS API is not supported.

Topics

- Asynchronous APIs in JavaScript
- Client-Side Slice-to-JavaScript Mapping
- Server-Side Slice-to-JavaScript Mapping

Copyright 2016, ZeroC, Inc.

Copyright 2016, ZeroC, Inc.

Asynchronous APIs in JavaScript

Given JavaScript's lack of support for threads, Ice for JavaScript uses asynchronous semantics in every situation that has the potential to block, including both local and non-local invocations. Synchronous proxy invocations are not supported.

Here is an example of a simple proxy invocation:

JavaScript

```
var proxy = HelloPrx.uncheckedCast(...);
var asyncResult = proxy.sayHello();
asyncResult.then(
    function(ar) {
        // handle success...
    },
    function(ex, ar) {
        // handle failure...
    }
);
```

The API design is similar to that of other asynchronous Ice language mappings in that the return value of a proxy invocation is an instance of `Ice.AsyncResult`, through which the program configures callbacks to handle the eventual success or failure of the invocation. The JavaScript implementation of `Ice.AsyncResult` also provides promise functionality.

Certain operations of local Ice run-time objects can also block, either because their implementations make remote invocations, or because their purpose is to block until some condition is satisfied:

- `Communicator::destroy`
- `Communicator::waitForShutdown`
- `Communicator::createObjectAdapterWithRouter`
- `Communicator::flushBatchRequests`
- `Connection::close`
- `Connection::flushBatchRequests`
- `ObjectPrx::ice_getConnection`

These operations use the same asynchronous API as for proxy invocations. The example below shows how to call `ice_getConnection` on a proxy:

JavaScript

```
var communicator = ...;
var proxy = communicator.stringToProxy("...");
proxy.ice_getConnection().then(
    function(connection, asyncResult) {
        // got connection...
    },
    function(ex, asyncResult) {
        // connection failed...
    }
);
```


Copyright 2016, ZeroC, Inc.

Client-Side Slice-to-JavaScript Mapping

In this section, we present the client-side Slice-to-JavaScript mapping. The client-side Slice-to-JavaScript mapping defines how Slice data types are translated to JavaScript types, and how clients invoke operations, pass parameters, and handle errors. Much of the JavaScript mapping is intuitive. For example, Slice sequences map to JavaScript arrays, so there is essentially nothing new you have to learn in order to use Slice sequences in JavaScript.

Much of what appears in this chapter is reference material. We suggest that you skim the material on the initial reading and refer back to specific sections as needed. However, we recommend that you read at least the mappings for exceptions, interfaces, and operations in detail because these sections cover how to call operations from a client, pass parameters, and handle exceptions.

In order to use the JavaScript mapping, you should need no more than the Slice definition of your application and knowledge of the JavaScript mapping rules. In particular, looking through the generated code in order to discern how to use the JavaScript mapping is likely to be inefficient, due to the amount of detail. Of course, occasionally, you may want to refer to the generated code to confirm a detail of the mapping, but we recommend that you otherwise use the material presented here to see how to write your client-side code.

The `Ice` Scope
All of the APIs for the Ice run time are nested in the `Ice` scope, to avoid clashes with definitions for other libraries or applications. Some of the contents of the `Ice` scope are generated from Slice definitions; other parts of the `Ice` scope provide special-purpose definitions that do not have a corresponding Slice definition. We will incrementally cover the contents of the `Ice` scope throughout the remainder of the manual.

Topics

- JavaScript Mapping for Identifiers
- JavaScript Mapping for Modules
- JavaScript Mapping for Built-In Types
- JavaScript Mapping for Enumerations
- JavaScript Mapping for Structures
- JavaScript Mapping for Sequences
- JavaScript Mapping for Dictionaries
- JavaScript Mapping for Constants
- JavaScript Mapping for Exceptions
- JavaScript Mapping for Interfaces
- JavaScript Mapping for Operations
- JavaScript Mapping for Classes
- slice2js Command-Line Options

Copyright 2016, ZeroC, Inc.

JavaScript Mapping for Identifiers

A Slice identifier maps to an identical JavaScript identifier. For example, the Slice identifier `Clock` becomes the JavaScript identifier `Clock`. There is one exception to this rule: if a Slice identifier is the same as a JavaScript keyword or is an identifier reserved by the Ice run time (such as `checkedCast`), the corresponding JavaScript identifier is prefixed with an underscore. For example, the Slice identifier `while` is mapped as `_while`.

You should try to avoid such identifiers as much as possible.

A single Slice identifier often results in several JavaScript identifiers. For example, for a Slice interface named `Foo`, the generated JavaScript code uses the identifiers `Foo` and `FooPrx` (among others). If the interface has the name `while`, the generated identifiers are `_while` and `whilePrx` (*not* `_whilePrx`), that is, the underscore prefix is applied only to those generated identifiers that actually require it.

See Also

- Lexical Rules
- JavaScript Mapping for Modules
- JavaScript Mapping for Built-In Types
- JavaScript Mapping for Enumerations
- JavaScript Mapping for Structures
- JavaScript Mapping for Sequences
- JavaScript Mapping for Dictionaries
- JavaScript Mapping for Constants
- JavaScript Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

JavaScript Mapping for Modules

A Slice module maps to a JavaScript scope with the same name as the Slice module. The mapping preserves the nesting of the Slice definitions. For example:

Slice

```
// Definitions at global scope here...

module M1 {
    // Definitions for M1 here...
    module M2 {
        // Definitions for M2 here...
    };
};

// ...

module M1 {     // Reopen M1
    // More definitions for M1 here...
};
```

This definition maps to the corresponding JavaScript definitions:

JavaScript

```
(function(module, require, exports)
{
    var Ice = ...
    var __M = Ice.__M;
    var M1 = __M.module("M1");

    // Definitions for M1 here...

    M1.M2 = __M.module("M1.M2");

    // Definitions for M2 here...

    // More definitions for M1 here...

    exports.M1 = M1;
}
(typeof (global) !== ...);
```

The definitions exported by the generated Slice code for top-level modules add symbols to the global window object in browsers, or to the native exports object in NodeJS.

When developing for NodeJS, you include the various Ice components with `require` statements as shown below:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var Ice = require("icejs").Ice;
var IceGrid = require("icejs").IceGrid;
var IceStorm = require("icejs").IceStorm;
...
```

Importing the generated code for your own Slice definitions works the same way:

JavaScript

```
var M1 = require("M1Defs").M1;
```

This example assumes the generated Slice definitions are in the file `M1Defs.js`.

In a browser application, the necessary JavaScript files are usually loaded via HTML:

```
<script type="text/javascript" src="Ice.js"></script>
<script type="text/javascript" src="M1Defs.js"></script>
<script type="text/javascript">
    // You can now access Ice and M1 as global objects
</script>
```

The file `Ice.js` shown above only provides definitions for the Ice run time; you would need to explicitly include the corresponding files for any other Ice components that your application needs, such as `IceGrid.js`.

See Also

- Modules
- Using the Slice Compilers
- JavaScript Mapping for Identifiers
- JavaScript Mapping for Built-In Types
- JavaScript Mapping for Enumerations
- JavaScript Mapping for Structures
- JavaScript Mapping for Sequences
- JavaScript Mapping for Dictionaries
- JavaScript Mapping for Constants
- JavaScript Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

JavaScript Mapping for Built-In Types

On this page:

- Mapping of Slice Built-In Types to JavaScript Types
- JavaScript Mapping for Long Integers

Mapping of Slice Built-In Types to JavaScript Types

The Slice built-in types are mapped to JavaScript types as follows:

Slice	JavaScript
bool	`Boolean`
byte	`Number`
short	`Number`
int	`Number`
long	`Ice.Long`
float	`Number`
double	`Number`
string	`String`

Mapping of Slice built-in types to JavaScript.

JavaScript Mapping for Long Integers

JavaScript does not provide a type that is capable of fully representing a Slice `long` value (a 64-bit integer), therefore Ice provides the `Ice.Long` type. An instance of `Ice.Long` encapsulates the high and low order 32-bit words (in little endian format) representing the 64-bit `long` value and provides access to these values via the `high` and `low` properties. Instances also provide the `toNumber` method, which returns a `Number` if the long value can be correctly represented by the `Number` type's integer range, otherwise it returns `Number.POSITIVE_INFINITY` or `Number.NEGATIVE_INFINITY` for positive and negative values, respectively.

`Ice.Long` objects have limited functionality because their primary purpose is to allow a `long` value to be re-transmitted, but instances do support the methods `hashCode`, `equals`, and `toString`.

See Also

- Basic Types
- JavaScript Mapping for Identifiers
- JavaScript Mapping for Modules
- JavaScript Mapping for Enumerations
- JavaScript Mapping for Structures
- JavaScript Mapping for Sequences
- JavaScript Mapping for Dictionaries
- JavaScript Mapping for Constants
- JavaScript Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

JavaScript Mapping for Enumerations

JavaScript does not have an enumerated type, so a Slice enumeration is emulated using JavaScript objects where each enumerator is an instance of the same type. For example:

Slice

```
enum Fruit { Apple, Pear, Orange };
```

The generated JavaScript code is equivalent to the following:

JavaScript

```
var Fruit = function(name, value) { ... };
Fruit.Apple = new Fruit("Apple", 0);
Fruit.Pear = new Fruit("Pear", 1);
Fruit.Orange = new Fruit("Orange", 2);
```

Each enumerator defines `name` and `value` properties that supply the enumerator's name and ordinal value, respectively. Enumerators also define `hashCode`, `equals` and `toString` methods, and the enumerated type itself defines a `valueOf` method that converts ordinal values into their corresponding enumerators.

Suppose we modify the Slice definition to include a custom enumerator value:

Slice

```
enum Fruit { Apple, Pear = 3, Orange };
```

The generated code changes accordingly:

JavaScript

```
var Fruit = function(name, value) { ... };
Fruit.Apple = new Fruit("Apple", 0);
Fruit.Pear = new Fruit("Pear", 3);
Fruit.Orange = new Fruit("Orange", 4);
```

Given the above definitions, we can use enumerated values as follows:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var f1 = Fruit.Apple;
var f2 = Fruit.Orange;

if (f1 === Fruit.Apple)   // Compare with constant
    // ...

if (f1 === f2)            // Compare two enums
    // ...

switch (f2) {             // Switch on enum
case Fruit.Apple:
    // ...
    break;
case Fruit.Pear:
    // ...
    break;
case Fruit.Orange:
    // ...
    break;
}

var f = Fruit.valueOf(3); // Convert an ordinal value to its enumerator, or
undefined if no match
console.log(f.name + " = " + f.value); // Outputs "Pear = 3"
```

Note that the generated type may contain other members, which we have not shown. These members are internal to the Ice run time and you must not use them in your application code (because they may change from release to release).

See Also

- Enumerations
- JavaScript Mapping for Structures
- JavaScript Mapping for Sequences
- JavaScript Mapping for Dictionaries

Copyright 2016, ZeroC, Inc.

JavaScript Mapping for Structures

On this page:

- Basic JavaScript Mapping for Structures
- JavaScript Constructors for Structures

Basic JavaScript Mapping for Structures

A Slice structure maps to a JavaScript type with the same name. For each Slice data member, the JavaScript instance contains a corresponding property. As an example, here is our Employee structure once more:

Slice

```
struct Employee {
    long number;
    string firstName;
    string lastName;
};
```

The JavaScript mapping for this structure is equivalent to the following code:

JavaScript

```
var Employee = function(number, firstName, lastName)
{
    this.number = ...
    this.firstName = ...
    this.lastName = ...
};
```

For each data member in the Slice definition, the JavaScript constructor assigns a value to its corresponding property. Instances define an `equals` method for comparison purposes and a `clone` method to create a shallow copy.

If the structure qualifies for use as a dictionary key, instances also define a `hashCode` function as required by the `Ice.HashMap` type.

JavaScript Constructors for Structures

The generated constructor has one parameter for each data member. This allows you to construct and initialize an instance in a single statement (instead of first having to construct the instance and then assign to its members). The constructor assigns a default value to each property appropriate for the type of its corresponding member:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False

Copyright 2016, ZeroC, Inc.

`sequence`	Null
`dictionary`	Null
`class/interface`	Null

If you wish to ensure that data members of primitive and enumerated types are initialized to specific values, you can declare default values in your Slice definition. The constructor initializes each of these data members to its declared value instead.

See Also

- Structures
- JavaScript Mapping for Enumerations
- JavaScript Mapping for Sequences
- JavaScript Mapping for Dictionaries

Copyright 2016, ZeroC, Inc.

JavaScript Mapping for Sequences

On this page:

- Basic JavaScript Mapping for Sequences
- JavaScript Mapping for Byte Sequences

Basic JavaScript Mapping for Sequences

A Slice sequence maps to a native JavaScript array. This means that the Slice-to-JavaScript compiler does not generate a separate named type for a Slice sequence. It also means that you can take advantage of all the inherent functionality offered by JavaScript's array type. For example:

Slice

```
sequence<Fruit> FruitPlatter;
```

We can use the `FruitPlatter` sequence as shown below:

JavaScript

```
var platter = [Fruit.Apple];
platter.push(Fruit.Pear);
```

JavaScript Mapping for Byte Sequences

As an optimization, occurrences of `sequence<byte>` map to a native buffer type, which is more efficient than native arrays for working with binary data. For browser applications, `sequence<byte>` maps to `Uint8Array`, and for Node.js applications it maps to the native `Buffer` type.

See Also

- Sequences
- JavaScript Mapping for Enumerations
- JavaScript Mapping for Structures
- JavaScript Mapping for Dictionaries

Copyright 2016, ZeroC, Inc.

JavaScript Mapping for Dictionaries

Here is the definition of our EmployeeMap once more:

Slice

```
dictionary<long, Employee> EmployeeMap;
```

In the JavaScript mapping, all dictionaries are instances of the `Ice.HashMap` type. The mapping also generates a type-specific constructor with the same name as the dictionary. For example:

JavaScript

```
var em = new EmployeeMap();

var e = new Employee();
e.number = new Ice.Long(31);
e.firstName = "James";
e.lastName = "Gosling";

em.set(e.number, e);
```

`HashMap` objects support the following properties and functions:

- `HashMap(keyComparator, valueComparator)`
  This version of the constructor accepts optional comparator functions that the map uses to compare keys and values for equality. If you instantiate a map directly using `new Ice.HashMap()` without specifying comparator functions, the default comparators use the `===` operator to compare keys and values. As an example, the following map compares its keys and values using `equals` meth ods:

  JavaScript

  ```
  function compareEquals(a, b) { return a.equals(b); }
  var m = new Ice.HashMap(compareEquals, compareEquals);
  ```

  The `valueComparator` function is only used when comparing two maps for equality.

  > The type-specific constructor generated for a Slice dictionary supplies comparator functions appropriate for its key and value types.

- `HashMap(h)`
  This version of the constructor accepts an optional `HashMap` object from which all entries are (shallow) copied. The new map inherits the comparator functions of the original.
- `size`
  The `size` property indicates how many entries are currently in the map.
- `entries`
  The `entries` property holds the head of a linked list that can be used to traverse all map entries.
- `get(key)`

Copyright 2016, ZeroC, Inc.

Returns the value associated with the given key, or `undefined` if no entry has a matching key.

- `set(key, value)`
  Adds a new entry to the map having the given key and value. Always returns `undefined`.
- `has(key)`
  Returns true if the map contains an entry with a matching key, or false otherwise.
- `delete(key)`
  Removes the entry associated with the given key. Returns the entry's current value if a match was found, or `undefined` otherwise.
- `clear()`
  Removes all entries in the current map.
- `keys()`
  Returns an array containing all keys in an unspecified order.
- `values()`
  Returns an array containing all values in an unspecified order.
- `forEach(fn, thisArg)`
  Iterates over the entries in the map in an unspecified order, calling the given function once for each entry. The optional thisArg parameter is used as `this` when executing the function. The arguments to the function are `key` and `value`, representing the key and value of each element.
- `equals(other, valueComparator)`
  Returns true if this map compares equal to the given map, false otherwise. You can optionally supply a function for `valueComparator` that the map uses when comparing values; this function takes precedence over the comparator supplied to the map's constructor.
- `clone()`
  Returns a shallow copy of the map.

Legal key types for `HashMap` include JavaScript's primitive types along with `null`, `NaN`, and any object that defines a `hashCode` method. The generated code for a Slice structure that qualifies as a legal dictionary key type includes a `hashCode` method. Suppose we define another dictionary type:

Slice

```
dictionary<Employee, string> EmployeeDeptMap;
```

The Slice compiler generates a constructor function equivalent to the following code:

JavaScript

```
var EmployeeDeptMap = function(h)
{
    var keyComparator = ...;
    var valueComparator = ...;
    return new Ice.HashMap(h || keyComparator, valueComparator);
};
```

Since the key is a user-defined structure type, the map requires a comparator that properly compares keys. Instantiating a map using `new EmployeeDeptMap` automatically sets the comparators, whereas calling `new Ice.HashMap` in this case would require you to supply your own comparators.

Ice applications should always instantiate a Slice dictionary using the generated constructor.

Copyright 2016, ZeroC, Inc.

It is not possible to use JavaScript's native `for..in` syntax to iterate over the map's entries. In addition to the `forEach` method described above, the map also maintains an unordered linked list of its entries that you can iterate through as shown below:

JavaScript

```
var em = new EmployeeMap();
// ...
for(var e = em.entries; e !== null; e = e.next)
{
    var employeeNum = e.key;
    var employee = e.value;
    ...
}
```

Each entry object supports the read-only properties `key`, `value` and `next`.

Notes

- Attempting to use the `map[key] = value` syntax to add an element to the map will not have the desired effect; you must use the `set` function instead.
- The Ice run time validates the elements of a dictionary to ensure that they are compatible with the declared type; an `Error` exceptio n is raised if an incompatible type is encountered.
- The `HashMap` API is largely compatible with the draft EcmaScript 6 `Map` type, but there are also some differences between the two. Ice for JavaScript may replace `HashMap` with the standard `Map` when it is widely supported.

See Also

- Dictionaries
- JavaScript Mapping for Enumerations
- JavaScript Mapping for Structures
- JavaScript Mapping for Sequences

Copyright 2016, ZeroC, Inc.

JavaScript Mapping for Constants

Here are the sample constant definitions once more:

Slice

```
module Example {
const bool      AppendByDefault = true;
const byte      LowerNibble = 0x0f;
const string    Advice = "Don't Panic!";
const short     TheAnswer = 42;
const double    PI = 3.1416;

enum Fruit { Apple, Pear, Orange };
const Fruit     FavoriteFruit = Pear;
};
```

For each constant, the JavaScript mapping generates a read-only property of the same name in the enclosing scope as shown below:

JavaScript

```
Object.defineProperty(Example, 'AppendByDefault', {value: true});
Object.defineProperty(Example, 'LowerNibble', {value: 15});
Object.defineProperty(Example, 'Advice', {value: "Don't Panic!"});
Object.defineProperty(Example, 'TheAnswer', {value: 42});
Object.defineProperty(Example, 'PI', {value: 3.1416});
Object.defineProperty(Example, 'FavoriteFruit', {value: Fruit.Pear});
```

Slice string literals that contain non-ASCII characters or universal character names are mapped to JavaScript string literals with universal character names. For example:

Slice

```
const string Egg = "œuf";
const string Heart = "c\u0153ur";
const string Banana = "\U0001F34C";
```

is mapped to:

JavaScript

```
Object.defineProperty(Example, 'Egg', {value: "\u0153uf"});
Object.defineProperty(Example 'Heart', {value: "c\u0153ur"});
Object.defineProperty(Example, 'Banana', {value: "\ud83c\udf4c"});
```

See Also

- Constants and Literals

Copyright 2016, ZeroC, Inc.

- JavaScript Mapping for Identifiers
- JavaScript Mapping for Modules
- JavaScript Mapping for Built-In Types
- JavaScript Mapping for Enumerations
- JavaScript Mapping for Structures
- JavaScript Mapping for Sequences
- JavaScript Mapping for Dictionaries
- JavaScript Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

JavaScript Mapping for Exceptions

On this page:

- JavaScript Mapping for User Exceptions
- JavaScript Constructors for User Exceptions
- JavaScript Mapping for Run-Time Exceptions

JavaScript Mapping for User Exceptions

Here is a fragment of the Slice definition for our world time server once more:

Slice

```
exception GenericError {
    string reason;
};
exception BadTimeVal extends GenericError {};
exception BadZoneName extends GenericError {};
```

These exception definitions map as follows:

JavaScript

```
GenericError = function(reason, cause) { ... }
GenericError.prototype = new Ice.UserException();
GenericError.prototype.constructor = GenericError;
GenericError.prototype.ice_name = function() { return "M::GenericError"; }

BadTimeVal = function(reason, cause) { ... }
BadTimeVal.prototype = new GenericError();
BadTimeVal.prototype.constructor = BadTimeVal;
BadTimeVal.prototype.ice_name = function() { return "M::BadTimeVal"; }

BadZoneName = function(reason, cause) { ... }
BadZoneName.prototype = new GenericError();
BadZoneName.prototype.constructor = BadZoneName;
BadZoneName.prototype.ice_name = function() { return "M::BadZoneName"; }
```

Each Slice exception maps to a JavaScript type with the same name. Each data member of an exception corresponds to a property in the JavaScript type.

The inheritance structure of the Slice exceptions is preserved for the generated types, so the prototypes for `BadTimeVal` and `BadZoneName` inherit from `GenericError`.

Each exception also defines the `ice_name` member function, which returns the name of the exception.

All user exceptions ultimately derive from the base type `Ice.UserException`. This allows you to handle any user exception generically as an `Ice.UserException`. Similarly, you can handle any Ice run-time exceptions as an `Ice.LocalException`, and you can handle any Ice exception as an `Ice.Exception`.

Note that the generated exception types may contain other properties and functions that are not shown. However, these elements are internal to the JavaScript mapping and are not meant to be used by application code.

Copyright 2016, ZeroC, Inc.

JavaScript Constructors for User Exceptions

The generated constructor has one parameter for each data member. This allows you to construct and initialize an instance in a single statement (instead of first having to construct the instance and then assign to its members). For derived exceptions, the constructor accepts one argument for each base exception member, plus one argument for each derived exception member, in base-to-derived order.

The constructor assigns a default value to each property appropriate for the type of its corresponding member:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	Null
`dictionary`	Null
`class/interface`	Null

If you wish to ensure that data members of primitive and enumerated types are initialized to specific values, you can declare default values in your Slice definition. The constructor initializes each of these data members to its declared value instead.

The generated constructor accepts an optional trailing argument representing the entity that caused the exception. This value is typically an instance of `Error`, but that is not a requirement. The value is used to initialize the `ice_cause` property inherited from `Ice.Exception`.

JavaScript Mapping for Run-Time Exceptions

The Ice run time throws run-time exceptions for a number of pre-defined error conditions. All run-time exceptions directly or indirectly derive from `Ice.LocalException` (which, in turn, derives from `Ice.Exception`).

Recall the inheritance diagram for user and run-time exceptions. By catching exceptions at the appropriate point in the hierarchy, you can handle exceptions according to the category of error they indicate:

- `Ice.Exception`
  This is the root of the inheritance tree for both run-time and user exceptions.
- `Ice.LocalException`
  This is the root of the inheritance tree for run-time exceptions.
- `Ice.UserException`
  This is the root of the inheritance tree for user exceptions.
- `Ice.TimeoutException`
  This is the base exception for both operation-invocation and connection-establishment timeouts.
- `Ice.ConnectTimeoutException`
  This exception is raised when the initial attempt to establish a connection to a server times out.

For example, a `ConnectTimeoutException` can be handled as `ConnectTimeoutException`, `TimeoutException`, `LocalException`, or `Exception`.

You will probably have little need to catch run-time exceptions as their most-derived type and instead catch them as `LocalException`; the fine-grained error handling offered by the remainder of the hierarchy is of interest mainly in the implementation of the Ice run time. Exceptions to this rule are the exceptions related to facet and object life cycles, which you may want to catch explicitly. These exceptions are `FacetNotExistException` and `ObjectNotExistException`, respectively.

See Also

- User Exceptions
- Run-Time Exceptions
- JavaScript Mapping for Modules
- JavaScript Mapping for Built-In Types

Copyright 2016, ZeroC, Inc.

- JavaScript Mapping for Enumerations
- JavaScript Mapping for Structures
- JavaScript Mapping for Sequences
- JavaScript Mapping for Dictionaries
- JavaScript Mapping for Constants
- Facets and Versioning
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

JavaScript Mapping for Interfaces

The mapping of Slice interfaces revolves around the idea that, to invoke a remote operation, you call a member function on a local class instance that is a proxy for the remote object. This makes the mapping easy and intuitive to use because making a remote procedure call is no different from making a local procedure call (apart from error semantics).

On this page:

- Proxy Interfaces in JavaScript
- The Ice.ObjectPrx Interface in JavaScript
- Down-casting Proxies in JavaScript
- Type IDs in JavaScript
- Using Proxy Methods in JavaScript
- Object Identity and Proxy Comparison in JavaScript

Proxy Interfaces in JavaScript

On the client side, a Slice interface maps to a JavaScript type with methods that correspond to the operations on that interface. Consider the following simple interface:

Slice

```
interface Simple {
    void op();
};
```

The Slice compiler generates code for use by the client similar to the following:

JavaScript

```
var SimplePrx = function() { ... }
SimplePrx.prototype.op = function(__ctx) { ... }
```

As you can see, the compiler generates a *proxy type* `SimplePrx`. In general, the generated name is `<`*`interface-name`*`>Prx`. If an interface is nested in a module `M`, the fully-qualified name is `M.<`*`interface-name`*`>Prx`.

In the client's address space, an instance of `SimplePrx` is the local ambassador for a remote instance of the `Simple` interface in a server and is known as a proxy instance. All the details about the server-side object, such as its address, what protocol to use, and its object identity are encapsulated in that instance.

Note that the prototype for `SimplePrx` inherits from `Ice.ObjectPrx`. This reflects the fact that all Ice interfaces implicitly inherit from `Ice::Object`.

For each operation in the interface, the proxy type defines a function with the same name. For the preceding example, we find that the operation `op` has been mapped to the function `op`. The function has an optional parameter `__ctx` that is used by the Ice run time to store information about how to deliver a request. You normally do not need to use it. Legal values for this parameter are `null` or an instance of `Ice.Context` (which is an `Ice.HashMap`). (We examine the `__ctx` parameter in detail in Request Contexts. The parameter is also used by IceStorm.)

> **Asynchronous Operations**
> Unlike most other Ice language mappings, where a Slice operation generates both synchronous and asynchronous proxy methods, the JavaScript language mapping generates only an asynchronous method. Given the event-driven nature of JavaScript and its networking API, it is not feasible to provide the traditional blocking RPC model.

Client code must not create an instance of a `<`*`interface-name`*`>Prx` type directly. Instead, proxy instances are always created on behalf of the client by the Ice run time, so client code never has any need to instantiate a proxy directly.

Copyright 2016, ZeroC, Inc.

A value of `null` denotes the null proxy. The null proxy is a dedicated value that indicates that a proxy points "nowhere" (denotes no object).

The `Ice.ObjectPrx` Interface in JavaScript

All Ice objects have `Object` as the ultimate ancestor type, so all proxies inherit from `Ice.ObjectPrx`. `ObjectPrx` provides a number of functions:

JavaScript

```
ObjectPrx.prototype.equals = function(r) { ... }
ObjectPrx.prototype.ice_getIdentity = function() { ... }
ObjectPrx.prototype.ice_isA = function(id, __ctx) { ... }
ObjectPrx.prototype.ice_ids = function(__ctx) { ... }
ObjectPrx.prototype.ice_id = function(__ctx) { ... }
ObjectPrx.prototype.ice_ping = function(__ctx) { ... }
```

The functions behave as follows:

- **`equals`**
  This function compares two proxies for equality. Note that all aspects of proxies are compared by this function, such as the communication endpoints for the proxy. This means that, in general, if two proxies compare unequal, that does *not* imply that they denote different objects. For example, if two proxies denote the same Ice object via different transport endpoints, `equals` returns `false` even though the proxies denote the same object.

- **`ice_getIdentity`**
  This function returns the identity of the object denoted by the proxy. The identity of an Ice object has the following Slice type:

  Slice

  ```
  module Ice {
  struct Identity {
      string name;
      string category;
  };
  };
  ```

  To see whether two proxies denote the same object, first obtain the identity for each object and then compare the identities:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var prx1 = ...;
var prx2 = ...;
var ident1 = prx1.ice_getIdentity();
var ident2 = prx2.ice_getIdentity();

if (i1.equals(i2))
    // o1 and o2 denote the same object
else
    // o1 and o2 denote different objects
```

- **`ice_isA`**
  The `ice_isA` function determines whether the object supports a specific interface. The argument to `ice_isA` is a type ID. For example, to see whether a proxy of type `ObjectPrx` denotes a `Printer` object, we can write:

  JavaScript

  ```
  var prx = ...;
  if (prx !== null) {
      prx.ice_isA("::Printer").then(
          function(b)
          {
              if(b)
                  // Target object supports the Printer interface!
              else
                  // Oops, target object is a different type
          },
          function(ex)
          {
              // Handle failure
          });
  }
  ```

  Note that we are testing whether the proxy is null before attempting to invoke the `ice_isA` method. This avoids getting a run-time error if the proxy is null.

- **`ice_ids`**
  The `ice_ids` function obtains an array of strings representing all of the type IDs that the object supports.

- **`ice_id`**
  The `ice_id` function obtains the type ID of the object. Note that the type returned is the type of the actual object, which may be more derived than the static type of the proxy. For example, if we have a proxy of type `BasePrx`, with a static type ID of `::Base`, the actual type ID might be `::Base`, or it might something more derived, such as `::Derived`.

- **`ice_ping`**
  The `ice_ping` function provides a basic reachability test for the object. If the object can physically be contacted (that is, the object exists and its server is running and reachable), the call completes normally; otherwise, it throws an exception that indicates why the object could not be reached, such as `ObjectNotExistException` or `ConnectTimeoutException`.

As remote operations, the `ice_isA`, `ice_ids`, `ice_id`, and `ice_ping` functions support an optional trailing argument representing a request context. Also note that there are other methods in `ObjectPrx`, not shown here. These methods provide different ways to dispatch a call

Copyright 2016, ZeroC, Inc.

and also provide access to an object's facets.

Down-casting Proxies in JavaScript

In addition to the methods corresponding to Slice operations, the Slice-to-JavaScript compiler also generates two helper methods that support down-casting:

JavaScript

```
SimplePrx.checkedCast = function(prx, facet, __ctx) { ... }
SimplePrx.uncheckedCast = function(prx, facet) { ... }
```

For `checkedCast`, if the passed proxy is for an object of type `Simple`, or a proxy for an object with a type derived from `Simple`, the result of the cast is a non-null reference to a proxy of type `SimplePrx`; otherwise, if the passed proxy denotes an object of a different type (or if the passed proxy is null), the result of the cast is a null reference.

Note that a checked cast contacts the server. This is necessary because only the implementation of an object in the server has definite knowledge of the type of an object. Consequently, a checked cast is implemented as the asynchronous method `checkedCast`, which may result in a `ConnectTimeoutException` or an `ObjectNotExistException`.

Given a proxy of any type, you can use a checked cast to determine whether the corresponding object supports a given type, for example:

JavaScript

```
var prx = ...;  // Get a proxy from somewhere...
SimplePrx.checkedCast(prx).then(
    function(proxy)
    {
        if(proxy !== null)
            // Object supports the Simple interface...
        else
            // Object is not of type Simple...
    },
    function(ex)
    {
        // Handle failure
    });
```

In contrast, an unchecked cast does not contact the server and unconditionally returns a proxy of the requested type. However, if you do use `uncheckedCast`, you must be certain that the proxy really does support the type you are casting to; otherwise, if you get it wrong, you will most likely get a run-time exception when you invoke an operation on the proxy. The most likely error for such a type mismatch is `OperationNotExistException`. However, other exceptions, such as a marshaling exception are possible as well. And, if the object happens to have an operation with the correct name, but different parameter types, no exception may be reported at all and you simply end up sending the invocation to an object of the wrong type; that object may do rather nonsensical things. To illustrate this, consider the following two interfaces:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Process {
    void launch(int stackSize, int dataSize);
};

// ...

interface Rocket {
    void launch(float xCoord, float yCoord);
};
```

Suppose you expect to receive a proxy for a `Process` object and use an `uncheckedCast` to down-cast the proxy:

JavaScript

```
var obj = ...;                                // Get proxy...
var process = ProcessPrx.uncheckedCast(obj);  // No worries...
process.launch(40, 60);                       // Oops...
```

If the proxy you received actually denotes a `Rocket` object, the error will go undetected by the Ice run time: because `int` and `float` have the same size and because the Ice protocol does not tag data with its type on the wire, the implementation of `Rocket::launch` will simply misinterpret the passed integers as floating-point numbers.

In fairness, this example is somewhat contrived. For such a mistake to go unnoticed at run time, both objects must have an operation with the same name and, in addition, the run-time arguments passed to the operation must have a total marshaled size that matches the number of bytes that are expected by the unmarshaling code on the server side. In practice, this is extremely rare and an incorrect `uncheckedCast` typically results in a run-time exception.

Type IDs in JavaScript

You can discover the type ID string corresponding to an interface by calling the `ice_staticId` function on the proxy type:

JavaScript

```
var id = SimplePrx.ice_staticId();
```

As an example, for an interface `Simple` in module `M`, the `ice_staticId` function returns the string `"::M::Simple"`.

Using Proxy Methods in JavaScript

The base proxy class `ObjectPrx` supports a variety of methods for customizing a proxy. Since proxies are immutable, each of these "factory methods" returns a copy of the original proxy that contains the desired modification. For example, you can obtain a proxy configured with a ten second timeout as shown below:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var proxy = communicator.stringToProxy(...);
proxy = proxy.ice_timeout(10000);
```

A factory method returns a new proxy object if the requested modification differs from the current proxy, otherwise it returns the current proxy. With few exceptions, factory methods return a proxy of the same type as the current proxy, therefore it is generally not necessary to repeat a checked or unchecked cast after using a factory method. The only exceptions are the factory methods `ice_facet` and `ice_identity`. Calls to either of these methods may produce a proxy for an object of an unrelated type, therefore they return a base proxy that you must subsequently down-cast to an appropriate type.

Object Identity and Proxy Comparison in JavaScript

Proxies provide an `equals` function that compares proxies:

JavaScript

```
ObjectPrx.prototype.equals = function(other) { ... }
```

Note that proxy comparison with `equals` uses *all* of the information in a proxy for the comparison. This means that not only the object identity must match for a comparison to succeed, but other details inside the proxy, such as the protocol and endpoint information, must be the same. In other words, comparison with `equals` tests for *proxy* identity, *not* object identity. A common mistake is to write code along the following lines:

JavaScript

```
var p1 = ...;        // Get a proxy...
var p2 = ...;        // Get another proxy...

if(!p1.equals(p2)) {
    // p1 and p2 denote different objects       // WRONG!
} else {
    // p1 and p2 denote the same object         // Correct
}
```

Even though `p1` and `p2` differ, they may denote the same Ice object. This can happen because, for example, both `p1` and `p2` embed the same object identity, but each use a different protocol to contact the target object. Similarly, the protocols may be the same, but denote different endpoints (because a single Ice object can be contacted via several different transport endpoints). In other words, if two proxies compare equal with `equals`, we know that the two proxies denote the same object (because they are identical in all respects); however, if two proxies compare unequal with `equals`, we know absolutely nothing: the proxies may or may not denote the same object.

To compare the object identities of two proxies, you can use a helper function:

Copyright 2016, ZeroC, Inc.

JavaScript

```
Ice.proxyIdentityCompare = function(lhs, rhs) { ... }
Ice.proxyIdentityAndFacetCompare = function(lhs, rhs) { ... }
```

`proxyIdentityCompare` allows you to correctly compare proxies for identity:

JavaScript

```
var p1 = ...;        // Get a proxy...
var p2 = ...;        // Get another proxy...

if (Ice.proxyIdentityCompare(p1, p2) !== 0) {
    // p1 and p2 denote different objects       // Correct
} else {
    // p1 and p2 denote the same object         // Correct
}
```

The function returns 0 if the identities are equal, -1 if `p1` is less than `p2`, and 1 if `p1` is greater than `p2`. (The comparison uses `name` as the major and `category` as the minor sort key.)

The `proxyIdentityAndFacetCompare` function behaves similarly, but compares both the identity and the facet name.

See Also

- Interfaces, Operations, and Exceptions
- Proxies for Ice Objects
- Type IDs
- JavaScript Mapping for Operations
- Request Contexts
- Operations on Object
- Proxy Methods
- Versioning

Copyright 2016, ZeroC, Inc.

JavaScript Mapping for Operations

JavaScript's event-driven APIs makes a synchronous invocation model impractical, therefore the JavaScript language mapping provides only an asynchronous model.

Asynchronous Method Invocation (AMI) is the term used to describe the client-side support for the asynchronous programming model. AMI supports both oneway and twoway requests and invocations never block. When a client issues an AMI request, the Ice run time hands the message off to the local transport buffer or, if the buffer is currently full, queues the request for later delivery. The application can then continue its activities and receive a callback when the invocation completes.

AMI is transparent to the server: there is no way for the server to tell whether a client sent a request synchronously or asynchronously.

On this page:

- Basic Asynchronous API in JavaScript
  - Asynchronous Proxy Methods in JavaScript
  - Passing Parameters in JavaScript
  - Null Parameters in JavaScript
  - Optional Parameters in JavaScript
  - Asynchronous Exception Semantics in JavaScript
- AsyncResult Type in JavaScript
- Promise Type in JavaScript
  - Basic Promise Concepts
  - Getting Started with Promises
  - Chaining Promises
  - Handling Errors with Promises
  - Passing Values between Promises
  - Chaining Asynchronous Events with Promises
  - Promise API
- Completion Callbacks in JavaScript
- Asynchronous Oneway Invocations in JavaScript
- Asynchronous Batch Requests in JavaScript

Basic Asynchronous API in JavaScript

Consider the following simple Slice definition:

Slice

```
module Demo {
    interface Employees {
        string getName(int number);
        string getAddress(int number);
    };
};
```

Asynchronous Proxy Methods in JavaScript

`slice2js` generates the following asynchronous proxy methods:

JavaScript

```
EmployeesPrx.prototype.getName = function(number, __ctx) { ... }
EmployeesPrx.prototype.getAddress = function(number, __ctx) { ... }
```


Copyright 2016, ZeroC, Inc.

As you can see, a Slice operation maps to a method of the same name. (Each function accepts an optional trailing per-invocation context.)

The `getName` function, for example, sends an invocation of `getName`. Because proxy invocations do not block, the application can do other things while the operation is in progress. The application receives notification about the completion of the request by registering callbacks for success and failure.

The `getName` function, as with all functions corresponding to Slice operations, returns a value of type `Ice.AsyncResult`. This value contains the state that the Ice run time requires to keep track of the asynchronous invocation.

A proxy function has one parameter for each in-parameter of the corresponding Slice operation. For example, consider the following operation:

Slice

```
double op(int inp1, string inp2, out bool outp1, out long outp2);
```

The `op` member function has the following signature:

JavaScript

```
function op(inp1, inp2, __ctx); // Returns Ice.AsyncResult
```

The operation's return value, as well as the values of its two out-parameters, are delivered to the success callback upon completion of the request.

Passing Parameters in JavaScript

The parameter passing rules for the JavaScript mapping are very simple: parameters are passed either by value (for simple types) or by reference (for complex types). Semantically, the two ways of passing parameters are identical: it is guaranteed that the value of a parameter will not be changed by the invocation.

Here is an interface with operations that pass parameters of various types from client to server:

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ClientToServer {
    void op1(int i, float f, bool b, string s);
    void op2(NumberAndString ns, StringSeq ss, StringTable st);
    void op3(ClientToServer* proxy);
};
```

The Slice compiler generates the following proxy methods for these definitions:

Copyright 2016, ZeroC, Inc.

JavaScript

```
ClientToServerPrx.prototype.op1 = function(i, f, b, s, __ctx) { ... }
ClientToServerPrx.prototype.op2 = function(ns, ss, st, __ctx) { ... }
ClientToServerPrx.prototype.op3 = function(proxy, __ctx) { ... }
```

Given a proxy to a `ClientToServer` interface, the client code can pass parameters as in the following example:

JavaScript

```
var p = ...;          // Get ClientToServerPrx proxy...

p.op1(42, 3.14, true, "Hello world!");  // Pass simple literals

var i = 42;
var f = 3.14;
var b = true;
var s = "Hello world!";
p.op1(i, f, b, s);                      // Pass simple variables

var ns = new NumberAndString();
ns.x = 42;
ns.str = "The Answer";
var ss = [];
ss.push("Hello world!");
var st = new StringTable();
st.set(0, ss);
p.op2(ns, ss, st);                      // Pass complex variables

p.op3(p);                               // Pass proxy
```

Null Parameters in JavaScript

Some Slice types naturally have "empty" or "not there" semantics. Specifically, sequences, dictionaries, and strings all can be `null`, but the corresponding Slice types do not have the concept of a null value. To make life with these types easier, whenever you pass `null` as a parameter or return value of type sequence, dictionary, or string, the Ice run time automatically sends an empty sequence, dictionary, or string to the receiver.

This behavior is useful as a convenience feature: especially for deeply-nested data types, members that are sequences, dictionaries, or strings automatically arrive as an empty value at the receiving end. This saves you having to explicitly initialize, for example, every string element in a large sequence before sending the sequence in order to avoid a run-time error. Note that using null parameters in this way does *not* create null semantics for Slice sequences, dictionaries, or strings. As far as the object model is concerned, these do not exist (only *empty* sequences, dictionaries, and strings do). For example, whether you send a string as `null` or as an empty string makes no difference to the receiver: either way, the receiver sees an empty string.

Optional Parameters in JavaScript

Optional parameters use the same mapping as required parameters. The only difference is that `undefined` can be passed as the value of an optional parameter or return value to indicate an "unset" condition. Consider the following operation:

Copyright 2016, ZeroC, Inc.

Slice

```
optional(1) int execute(optional(2) string params, out optional(3) float value);
```

A client can invoke this operation as shown below:

JavaScript

```
proxy.execute("--file log.txt").then(
    function(retval, value)
    {
        if(value !== undefined)
            console.log("value =", value);
    });
```

A well-behaved program must always compare an optional parameter to `undefined` prior to using its value.

For Slice types that support null semantics, such as proxies and objects by value, passing `null` for an optional parameter has a different meaning than passing `undefined`:

- `null` means a value was supplied for the parameter, and that value happens to be `null`
- `undefined` means no value was supplied for the parameter

`undefined` is not a legal value for required parameters.

Asynchronous Exception Semantics in JavaScript

If an invocation raises an exception, the exception is delivered to the failure callback, even if the actual error condition for the exception was encountered during the proxy function ("on the way out"). The advantage of this behavior is that all exception handling is located in the failure callback (instead of being present twice, once where the proxy function is called, and again in the failure callback).

There is one exception to the above rule: if you destroy the communicator and then make an asynchronous invocation, the proxy function throws `CommunicatorDestroyedException`. This is necessary because, once the run time is finalized, it can no longer deliver an exception to the failure callback.

The only other exception that is thrown by the proxy function is `Error`. This exception indicates that you have used the API incorrectly. For example, the proxy function throws this exception if you pass an invalid argument for one of the operation's parameters.

`AsyncResult` Type in JavaScript

The `AsyncResult` object returned by an Ice API call encapsulates the state of the asynchronous invocation.

In other Ice language mappings, the `AsyncResult` type is only used for asynchronous remote invocations. In JavaScript, an `AsyncResult` object can also be returned by methods of local Ice run time objects such as `Communicator` and `ObjectAdapter` when those methods have the potential to block or internally make remote invocations.

An instance of `AsyncResult` defines the following properties:

- `communicator`
  The communicator that sent the invocation.
- `connection`
  This method returns the connection that was used for the invocation. Note that, for typical proxy invocations, this method returns a nil value because the possibility of automatic retries means the connection that is currently in use could change unexpectedly. The c

Copyright 2016, ZeroC, Inc.

onnection property only returns a non-nil value when the AsyncResult object is obtained by calling flushBatchRequests on a Connection object.

- proxy
  The proxy that was used for the invocation, or nil if the AsyncResult object was not obtained via a proxy invocation.
- adapter
  The object adapter that was used for the invocation, or nil if the AsyncResult object was not obtained via an object adapter invocation.
- operation
  The name of the operation being invoked.

AsyncResult objects also support the following methods:

- cancel()
  This method prevents a queued invocation from being sent or, if the invocation has already been sent, ignores a reply if the server sends one. cancel is a local operation and has no effect on the server. A canceled invocation is considered to be completed, meaning isCompleted returns true, and the result of the invocation is an Ice.InvocationCanceledException.
- isCompleted()
  This method returns true if, at the time it is called, the result of an invocation is available, indicating that a call to the end_ method will not block the caller. Otherwise, if the result is not yet available, the method returns false.
- isSent()
  When you make a proxy invocation, the Ice run time attempts to write the corresponding request to the client-side transport. If the transport cannot accept the request, the Ice run time queues the request for later transmission. isSent returns true if, at the time it is called, the request has been written to the local transport (whether it was initially queued or not). Otherwise, if the request is still queued or an exception occurred before the request could be sent, isSent returns false.
- throwLocalException()
  This method throws the local exception that caused the invocation to fail. If no exception has occurred yet, throwLocalException does nothing.
- sentSynchronously()
  This method returns true if a request was written to the client-side transport without first being queued. If the request was initially queued, sentSynchronously returns false (independent of whether the request is still in the queue or has since been written to the client-side transport).

As described in the next section, an AsyncResult object provides additional functionality, including support for success and failure callbacks, that it inherits from a base type.

Promise Type in JavaScript

The Ice.AsyncResult type actually derives from Ice.Promise. This type is an implementation of the *promise* concept commonly used in JavaScript applications for simplifying the handling of asynchronous invocations. As you'll see below, one especially useful feature of a promise is its support for chaining, which allows you to make a series of asynchronous calls while still maintaining code readability.

Basic Promise Concepts

A promise is a relatively simple object whose purpose is to invoke callback functions upon the success or failure of a condition. The object always starts out in a pending state, and then transitions to a success or failure state via calls to its succeed and fail methods, respectively. This state transition, known as *resolving* a promise, occurs only once, with any subsequent calls to succeed or fail ignored. The initial caller of succeed or fail can supply arguments that the promise object will pass along to its success or failure callbacks. Any number of callbacks can be registered on a single promise, and they can be registered before or after the promise has been resolved.

The promise implementation always invokes its callbacks asynchronously via setTimeout, which avoids the complexities that could occur if the callbacks were invoked immediately as nested calls. For the purposes of this discussion, when we say that a promise "invokes" its callbacks, we actually mean that it queues the callbacks for invocation.

A promise invokes any registered callbacks at the time the promise is resolved. Callbacks registered after a promise has already been resolved are invoked immediately. The promise keeps references to the arguments passed to succeed or fail so that it can pass them to callbacks registered later.

Getting Started with Promises

Copyright 2016, ZeroC, Inc.

The `Promise` method you'll use most often is `then`, which accepts two arguments:

JavaScript

```
function then(successFn, failureFn) { ... }
```

Both arguments are optional. Here is a simple example:

JavaScript

```
var p = new Ice.Promise();
p.then(
    function(str, i) {
        console.log("received " + str + " and " + i);
    },
    function(ex) {
        console.log("failed: " + ex);
    });

try {
    // do something...
    p.succeed("string arg", 5);
}
catch(ex) {
    p.fail(ex);
}
```

The call to `succeed` resolves the promise and supplies two arguments that the promise passes to the success callback. Similarly, the call to `fail` passes along the exception caught by the caller.

As you can see, the success and fail callbacks typically must know what arguments to expect when the promise is resolved.

Chaining Promises

The convenience and power offered by the promise concept becomes clear when you need to execute a sequence of asynchronous actions. Consider this example:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var p = new Ice.Promise();
p.then(    // A
    function() {
        // Step 1...
    }
).then(    // B
    function() {
        // Step 2...
    }
).then(    // C
    function() {
        // Step 3...
    }
);

// None of this happens until...
p.succeed();
```

Each call to `then` returns a new promise that is automatically chained to the preceding one. The promise returned by `then` in A resolves when Step 1 completes. The successful completion of A causes Step 2 to execute, and the successful completion of Step 2 resolves the promise returned by `then` in B, triggering Step 3 to execute. Finally, the whole chain of events won't begin until the initial promise `p` resolves.

Handling Errors with Promises

Let's extend our previous example to include an error handler:

JavaScript

```
var p = new Ice.Promise();
p.then(    // A
    function() {
        // Step 1...
    }
).then(    // B
    function() {
        // Step 2...
    }
).exception(
    function(ex) {
        // handle errors
    }
);
```

Here we call the `exception(failureFn)` method to establish an exception handler at the end of the promise chain, which is equivalent to calling `then(undefined, failureFn)`. You'll notice that none of the intermediate steps defines its own failure callback, with the intention that all errors will be handled in one location. After a call to `p.succeed` triggers execution of the chain, an exception raised by any of these steps propagates through each subsequent promise in the chain until a failure callback is found. If instead the application calls `p.fail`, none of the steps would execute and only the exception handler would be called.

Copyright 2016, ZeroC, Inc.

Now let's examine the behavior when an intermediate step defines a failure callback:

JavaScript

```
var p = new Ice.Promise();
p.then(
    function() {
        console.log("step 1");
    },
    function() {
        console.log("fail 1");
    }
).then(
    function() {
        console.log("step 2");
    }
).exception(
    function(ex) {
        console.log("error");
    }
);
p.fail();

// Outputs:
fail 1
step 2
```

The results might surprise you at first. If a failure callback completes without raising an exception, the promise is considered to have resolved successfully and therefore triggers the success callback of the next promise in the chain. On the other hand, if a failure callback raises an exception, it will propagate to the next failure callback in the chain. This behavior allows a failure callback to recover from error situations (if possible) and continue with the normal logic flow, or raise an exception and skip to the next error handler.

If no failure callback is defined in a chain after the point at which an exception occurs, that exception will be silently ignored. For this reason, we always recommend terminating a chain with an exception handler. Furthermore, you should not let any uncaught exceptions be thrown from your final exception handler.

Next we'll illustrate this point by throwing an exception:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var p = new Ice.Promise();
p.then(
    function() {
        console.log("step 1");
    },
    function() {
        console.log("fail 1");
        throw new Error();
    }
).then(
    function() {
        console.log("step 2");
    }
).exception(
    function(ex) {
        console.log("error");
    }
);
p.fail();

// Outputs:
fail 1
error
```

Here you can see that the exception raised in the first failure callback bypasses Step 2 and triggers the final exception handler.

Passing Values between Promises

The promise implementation forwards any value returned by a success or failure callback to the next success callback in the chain, as shown below:

JavaScript

```
var p = new Ice.Promise();
p.then(
    function() {
        return "World!";
    }
).then(
    function(arg) {
        console.log("Hello " + arg);
    }
);
p.succeed();

// Outputs:
Hello World!
```


Copyright 2016, ZeroC, Inc.

Chaining Asynchronous Events with Promises

Now that we've discussed promise fundamentals, let's explore more realistic use cases. The primary purpose of promises is chaining together a sequence of asynchronous events such that the next step in the chain won't execute until the previous step has completed. Common examples in Ice applications include a series of proxy invocations in which the calls must be made sequentially in a certain order, and proxy invocations in which the results of a preceding call must be obtained before the next call can be made.

Consider the following example:

JavaScript

```
var p = new Ice.Promise();
p.then(
    function() {  // Step 1
        var promise = // do something that creates a promise...
        return promise;
    }
).then(
    function() {  // Step 2
        var promise = // do something else that creates a promise...
        return promise;
    }
).then(
    function() {  // Step 3
        console.log("all done");
    }
);

p.succeed();
```

Steps 1 and 2 each initiate some asynchronous action that creates a promise, and the steps return the promise object as the result of the callback. Normally the result value would be passed as the argument to the next success callback in the chain, as we discussed earlier. However, when a callback returns a promise, execution of the next success or failure callback becomes dependent on the completion of the returned promise. In the example above, calling `p.succeed` causes Step 1 to execute, but Step 2 won't execute until the promise returned by Step 1 resolves successfully.

Promise API

Instances of `Ice.Promise` support the following methods:

- `then(successFn, failureFn)`
  Executes the success function if this promise resolves successfully, or the failure function if this promise fails. `then` returns a new promise whose execution is dependent on the resolution of this promise. If either function argument is undefined, the success or failure status of this promise is passed to the next one in the chain. If a function is undefined and this is the last promise in the chain, the outcome is ignored.
- `exception(failureFn)`
  A convenience method equivalent to calling `then(undefined, failureFn)`. `exception` returns a new promise whose execution is dependent on the resolution of this promise.
- `finally(fn)`
  Executes the given function when this promise succeeds or fails. The callback function must be prepared to receive any arguments resulting from the success or failure of this promise. `finally` returns a new promise whose execution is dependent on the resolution of the previous promise in the chain (see example below). Exceptions raised by the given callback function will be ignored.
- `delay(ms)`
  Returns a new promise whose success or failure callback is invoked after the specified delay (in milliseconds). The results of the previous promise in the chain are forwarded to the next promise in the chain, as if the delay promise was not present.

Copyright 2016, ZeroC, Inc.

- `succeed([args])`
  Resolves this promise successfully. Any arguments are passed to the success callback of the next promise in the chain. Returns a reference to this promise.
- `fail([args])`
  Resolves this promise as failed. Any arguments are passed to the failure callback of the next promise in the chain. Returns a reference to this promise.
- `succeeded()`
  Returns true if this promise succeeded, false otherwise.
- `failed()`
  Returns true if this promise failed, false otherwise.
- `completed()`
  Returns true if this promise has been resolved, false otherwise.

Additionally, the `Promise` type defines the following functions:

- `try(fn)`
  Returns a new promise whose execution depends on the completion of the given function. This function allows you to write code similar to the traditional try/catch/finally blocks of synchronous code, as shown below:

JavaScript

```
Ice.Promise.try(
    function() {
        // perform asynchronous actions...

        // return a promise to the last action in the chain...
    }
).exception(
    function(ex) {
        // handle errors
    }
).finally(
    function() {
        // clean up when the last action completes
    }
);
```

  The `try` function makes error handling especially convenient because the chained exception handler will be called if an exception is thrown directly by the `try` callback, or if the callback returns a promise that later fails. Note that you could also arrange the `exception` and `finally` blocks this way:

Copyright 2016, ZeroC, Inc.

JavaScript

```
Ice.Promise.try(
    function() {
        // perform asynchronous actions...

        // return a promise to the last action in the chain...
    }
).finally(
    function() {
        // clean up when the last action completes
    }
).exception(
    function(ex) {
        // handle errors
    }
);
```

If the `try` callback fails with an exception, the `finally` callback will be invoked and the failure will propagate to the `exception` callback. Of course, the order in which the callbacks are invoked in an error situation differs from the previous example.

- `all(promise1[, promise2, ...])`
  Returns a new promise whose execution depends on the completion of all given promises. The promise arguments can be specified individually or supplied as a single array, meaning `all(p1, p2, p3)` is equivalent to `all([p1, p2, p3])`. The returned promise resolves successfully if all of the given promises resolve successfully, and the returned promise fails if any of the given promises fails. In the case that all of the given promises succeed, the result values from each promise's success function are passed as arguments to the success callback of the returned promise in the same order. Here is an example:

JavaScript

```
function asyncAction() {
    // returns a promise...
}

Ice.Promise.all(
    asyncAction(), // Action 1
    asyncAction(), // Action 2
    asyncAction()  // Action 3
).then(
    function(results1, results2, results3) {
        // Each argument is an array containing the results of the
corresponding action
    }
).exception(
    function(ex) {
        // handle errors
    }
);
```

The success function of the promise returned by `all` does not execute until all of the promise arguments resolve successfully, at which time their results are passed as arguments to the success function. Each result argument is an array containing the results (if

Copyright 2016, ZeroC, Inc.

any) of the corresponding action's promise.

- `delay([args...,] ms)`
  Similar to the prototype method described earlier, this function returns a promise whose success callback is invoked after the specified delay (in milliseconds). One difference between this function and the prototype method is the delay value must be the last argument, with any preceding arguments passed to the success callback:

JavaScript

```
Ice.Promise.delay("John", 100).then(
    function(name) { // Invoked after 100ms
        console.log("Hi " + name);
    },
    ...
);
```

Completion Callbacks in JavaScript

The `AsyncResult` promise returned by a proxy invocation resolves when the remote operation completes successfully or fails. The semantics of the success and failure functions depend on the proxy's invocation mode:

- Twoway proxies
  The success or failure function executes after the Ice run time in the client receives a reply from the server.
- Oneway and datagram proxies
  The success function executes after the Ice run time passes the message off to the socket. The failure function will only be called if an error occurs while Ice is attempting to marshal or send the message.
- Batch proxies
  The success function executes after the Ice run time queues the request. The failure function will only be called if an error occurs while Ice is attempting to marshal the message.

For all asynchronous Ice APIs, the failure function receives two arguments: the exception that caused the failure, and a reference to the `AsyncResult` promise for the failed invocation.

The success callback parameters depend on the operation signature. If the operation has non-`void` return type, the first parameter of the success callback is the return value. The return value (if any) is followed by a parameter for each out-parameter of the corresponding Slice operation, in the order of declaration. Ice adds a reference to the `AsyncResult` object associated with the request as the final parameter to every success callback. Recall the example we showed earlier:

Slice

```
double op(int inp1, string inp2, out bool outp1, out long outp2);
```

The success callback for `op` must have the following signature:

JavaScript

```
function handleSuccess(returnVal, outp1, outp2, asyncResult)
```

You can omit the `asyncResult` parameter if your function has no use for it.

The failure callback is invoked if the invocation fails because of an Ice run time or user exception. This callback function must have the

Copyright 2016, ZeroC, Inc.

following signature:

JavaScript

```
function handleAnException(ex, asyncResult)
```

Again, you can omit the `asyncResult` parameter if your function has no use for it.

For example, the code below shows how to invoke the `getName` operation:

JavaScript

```
var proxy = ...;
var empNo = ...;
proxy.getName(empNo).then(
    function(name, asyncResult) {
        console.log("Name of employee #" + empNo + " is " + name);
    }
).exception(
    function(ex, asyncResult) {
        console.log("Failed to get name of employee #" + empNo);
        console.log(ex);
    }
);
```

Let's extend the example add a call to `getAddress`:

JavaScript

```
var proxy = ...;
var empNo = ...;
proxy.getName(empNo).then(
    function(name, asyncResult) {
        console.log("Name of employee #" + empNo + " is " + name);
        return proxy.getAddress(empNo);
    }
).then(
    function(addr, asyncResult) {
        console.log("Address of employee #" + empNo + " is " + addr);
    }
).exception(
    function(ex, asyncResult) {
        console.log(asyncResult.operation + " failed for employee #" + empNo);
        console.log(ex);
    }
);
```

Notice here that the success callback for `getName` returns the `AsyncResult` promise for `getAddress`, which means the second success

Copyright 2016, ZeroC, Inc.

callback won't be invoked until `getAddress` completes.

Asynchronous Oneway Invocations in JavaScript

You can invoke operations via oneway proxies asynchronously, provided the operation has `void` return type, does not have any out-parameters, and does not raise user exceptions. If you call the proxy function on a oneway proxy for an operation that returns values or raises a user exception, the proxy function throws an instance of `Error`.

The callback functions look exactly as for a twoway invocation. The failure function is only called if the invocation raised an exception during the proxy function ("on the way out"), and the success function is called as soon as the message is accepted by the local network connection. The success function takes a single argument representing the `AsyncResult` associated with the request.

Asynchronous Batch Requests in JavaScript

Applications that send batched requests can either flush a batch explicitly or allow the Ice run time to flush automatically. The asynchronous proxy method `ice_flushBatchRequests` initiates an immediate flush of any batch requests queued by that proxy.

In addition, similar methods are available on the communicator and the `Connection` object that is accessible via the `connection` property of `AsyncResult`. These methods flush batch requests sent via the same communicator and via the same connection, respectively.

All versions of `ice_flushBatchRequests` return a promise whose success callback executes as soon as the batch request message is accepted by the local network connection. The success function takes a single argument representing the `AsyncResult` associated with the flush request. Ice only invokes the failure callback if an error occurs while attempting to send the message.

See Also

- Request Contexts
- Batched Invocations

Copyright 2016, ZeroC, Inc.

JavaScript Mapping for Classes

On this page:

- Basic JavaScript Mapping for Classes
- Inheritance from Ice.Object in JavaScript
- Class Constructors in JavaScript
- Class Data Members in JavaScript
- Class Operations in JavaScript
- Class Factories in JavaScript

Basic JavaScript Mapping for Classes

A Slice class is mapped to a JavaScript type with the same name. For each Slice data member, the JavaScript instance contains a corresponding property (just as for structures and exceptions). Consider the following class definition:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Slice compiler generates the following code for this definition:

JavaScript

```
TimeOfDay = function(hour, minute, second) { ... }
TimeOfDay.prototype = new Ice.Object();
TimeOfDay.prototype.constructor = TimeOfDay;
// Nothing generated for format()...
```

There are a number of things to note about the generated code:

1. The prototype for generated type `TimeOfDay` inherits from `Ice.Object`. This means that all Slice classes implicitly inherit from `Ice.Object`, which is the ultimate ancestor of all classes. Note that `Ice.Object` is *not* the same as `Ice.ObjectPrx`. In other words, you *cannot* pass a class where a proxy is expected and vice versa.
2. The generated type provides a constructor that accepts a value for each data member.
3. The generated type defines a property for each Slice data member.
4. The generated type's prototype does not include any definitions for the class' operations.

There is quite a bit to discuss here, so we will look at each item in turn.

Inheritance from `Ice.Object` in JavaScript

As for Slice interfaces, the generated type for a Slice class implicitly inherits from a common base type, `Ice.Object`. However, as shown in the illustration below, the types inherit from `Ice.Object` instead of `Ice.ObjectPrx` (which is at the base of the inheritance hierarchy for proxies). As a result, you cannot pass a class where a proxy is expected (and vice versa) because the base types for classes and proxies are not compatible.

Copyright 2016, ZeroC, Inc.

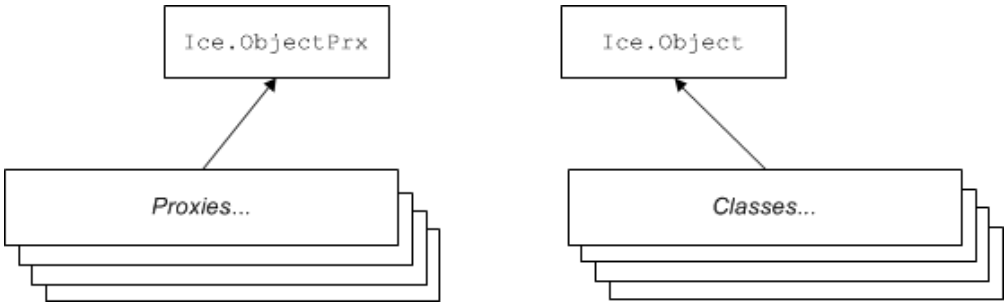

Inheritance from `Ice.ObjectPrx` and `Ice.Object`.

`Ice.Object` defines a number of member functions:

JavaScript

```
Object.prototype.ice_isA = function(s, current) { ... }
Object.prototype.ice_ping = function(current) { }
Object.prototype.ice_ids = function(current) { ... }
Object.prototype.ice_id = function(current) { ... }
Object.prototype.ice_preMarshal = function() { }
Object.prototype.ice_postMarshal = function() { }
```

The member functions of `Ice.Object` behave as follows:

- `ice_isA`
  This function returns `true` if the object supports the given type ID, and `false` otherwise.
- `ice_ping`
  As for interfaces, `ice_ping` provides a basic reachability test for the class.
- `ice_ids`
  This function returns a string sequence representing all of the type IDs supported by this object, including `::Ice::Object`.
- `ice_id`
  This function returns the actual run-time type ID for a class. If you call `ice_id` through a reference to a base instance, the returned type ID is the actual (possibly more derived) type ID of the instance.
- `ice_preMarshal`
  The Ice run time invokes this function prior to marshaling the object's state, providing the opportunity for a subtype to validate its declared data members.
- `ice_postUnmarshal`
  The Ice run time invokes this function after unmarshaling an object's state. A subtype typically overrides this function when it needs to perform additional initialization using the values of its declared data members.

Class Constructors in JavaScript

The type generated for a Slice class provides a constructor that initializes each data member to a default value appropriate for its type:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False

Copyright 2016, ZeroC, Inc.

`sequence`	Null
`dictionary`	Null
`class/interface`	Null

If you wish to ensure that data members of primitive and enumerated types are initialized to specific values, you can declare default values in your Slice definition. The constructor initializes each of these data members to its declared value instead.

The constructor accepts one argument for each member of the class. This allows you to create and initialize an instance in a single statement, for example:

JavaScript

```
var tod = new TimeOfDayI(14, 45, 00); // 14:45pm
```

For derived classes, the constructor requires an argument for every member of the class, including inherited members. For example, consider the the definition from Class Inheritance once more:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
};

class DateTime extends TimeOfDay {
    short day;          // 1 - 31
    short month;        // 1 - 12
    short year;         // 1753 onwards
};
```

The constructors generated for these classes are similar to the following:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var TimeOfDay = function(hour, minute, second)
{
    this.hour = hour;
    this.minute = minute;
    this.second = second;
};

var DateTime = function(hour, minute, second, day, month, year)
{
    TimeOfDay.call(this, hour, minute, second);
    this.day = day;
    this.month = month;
    this.year = year;
};
```

Pass `undefined` as the value of any optional data member that you wish to remain unset.

Class Data Members in JavaScript

By default, data members of classes are mapped exactly as for structures and exceptions: for each data member in the Slice definition, the generated type defines a corresponding property.

Optional data members use the same mapping as required data members, but an optional data member can also be set to `undefined` to indicate that the member is unset. A well-behaved program must compare an optional data member to `undefined` before using the member's value:

JavaScript

```
var v = ...
if(v.optionalMember === undefined)
    console.log("optionalMember is unset")
else
    console.log("optionalMember =", v.optionalMember)
```

Class Operations in JavaScript

Operations of classes are mapped to methods in the generated type. This means that, if a class contains operations (such as the `format` operation of our `TimeOfDay` class), you must provide an implementation of those operations in a type that is derived from the generated type. For example:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var TimeOfDayI = function(hour, minute, second)
{
    TimeOfDay.call(this, hour, minute, second);
};
TimeOfDayI.prototype = new TimeOfDay();
TimeOfDayI.prototype.constructor = TimeOfDayI;

TimeOfDayI.prototype.format = function(current)
{
    var result = "";
    result += (this.hour < 10) ? "0" + this.hour : this.hour;
    result += ":";
    result += (this.minute < 10) ? "0" + this.minute : this.minute;
    result += ":";
    result += (this.second < 10) ? "0" + this.second : this.second;
    return result;
};
```

Note that Ice provides an equivalent but much more compact way of defining class implementations:

JavaScript

```
var TimeOfDayI = Ice.Class(TimeOfDay, {
    format: function(current) { ... }
});
```

The `Ice.Class` constructor function creates a new type, establishes the prototype relationship (if a base type is specified), and adds each defined method to the new type's prototype.

Class Factories in JavaScript

Having created a type such as `TimeOfDayI`, we have an implementation and we can instantiate `TimeOfDayI` objects, but we cannot receive it as the return value or as an out-parameter from an operation invocation. To see why, consider the following simple interface:

Slice

```
interface Time {
    TimeOfDay get();
};
```

When a client invokes the `get` operation, the Ice run time must instantiate and return an instance of `TimeOfDay`. However, `TimeOfDay` is conceptually an abstract type because it declares operations and therefore it cannot be instantiated. Unless we tell it, the Ice run time cannot magically know that we have created a `TimeOfDayI` type that implements the abstract `format` operation of the `TimeOfDay` abstract type. In other words, we must provide the Ice run time with a factory that knows that the `TimeOfDay` abstract type has a `TimeOfDayI` concrete implementation. The `Ice::Communicator` interface provides us with the necessary operations:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local interface ObjectFactory {
        Object create(string type);
        void destroy();
    };

    local interface Communicator {
        void addObjectFactory(ObjectFactory factory, string id);
        ObjectFactory findObjectFactory(string id);
        // ...
    };
};
```

To supply the Ice run time with a factory for our `TimeOfDayI` type, we must define an object that implements the `create` and `destroy` met hods:

JavaScript

```
var factory = {
    create: function(type)
    {
        if(type === TimeOfDay.ice_staticId())
        {
            return new TimeOfDayI();
        }
        return null;
    },
    destroy: function()
    {
        // Nothing to do
    }
};
```

The factory's `create` function is called by the Ice run time when it needs to create a `TimeOfDay` instance. The factory's `destroy` function is called by the Ice run time when its communicator is destroyed.

The `create` function is passed the type ID of the class to instantiate. For our `TimeOfDay` class, the type ID is `"::M::TimeOfDay"`. Our implementation of `create` checks the type ID: if it matches, the function instantiates and returns a `TimeOfDayI` object. For other type IDs, the function returns null because it does not know how to instantiate other types of objects.

Note that we used the `ice_staticId` function to obtain the type ID rather than embedding a literal string. Using a literal type ID string in your code is discouraged because it can lead to errors that are only detected at run time. For example, if a Slice class or one of its enclosing modules is renamed and the literal string is not changed accordingly, a receiver will fail to unmarshal the object and the Ice run time will raise `NoObjectFactoryException`. By using `ice_staticId` instead, we avoid any risk of a misspelled or obsolete type ID, and we can more easily discover whether a Slice class or module has been renamed.

Given a factory implementation, such as the one above, we must inform the Ice run time of the existence of the factory:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var communicator = ...;
communicator.addObjectFactory(factory, TimeOfDay.ice_staticId());
```

Now, whenever the Ice run time needs to instantiate an object with the type ID `"::M::TimeOfDay"`, it calls the `create` function of the registered factory.

The `destroy` operation of the object factory is invoked by the Ice run time when the communicator is destroyed. This gives you a chance to clean up any resources that may be used by your factory. Do not call `destroy` on the factory while it is registered with the communicator — if you do, the Ice run time has no idea that this has happened and, depending on what your `destroy` implementation is doing, may cause undefined behavior when the Ice run time tries to next use the factory.

The run time guarantees that `destroy` will be the last call made on the factory, that is, `create` will not be called once `destroy` has been called.

Note that you cannot register a factory for the same type ID twice: if you call `addObjectFactory` with a type ID for which a factory is registered, the Ice run time throws an `AlreadyRegisteredException`.

Finally, keep in mind that if a Slice class has only data members, but no operations, you need not create and register an object factory to transmit instances of such a class. Only if a class has operations do you have to define and register an object factory.

See Also

- Classes
- Class Inheritance
- Type IDs

Copyright 2016, ZeroC, Inc.

slice2js Command-Line Options

The Slice-to-JavaScript compiler, `slice2js`, offers the following command-line options in addition to the standard options:

- `--stdout`
  Print generated code to standard output.
- `--depend-json`
  Print dependency information in JSON format to standard output by default, or to the file specified by the `--depend-file` option. N o code is generated when this option is specified. The output consists of the complete list of Slice files that the input Slice files depend on through direct or indirect inclusion.

See Also

- Using the Slice Compilers

Copyright 2016, ZeroC, Inc.

Server-Side Slice-to-JavaScript Mapping

The mapping for Slice data types to JavaScript is identical on the client side and server side. This means that everything in Client-Side Slice-to-JavaScript Mapping also applies to the server side. However, for the server side, there are a few additional things you need to know — specifically how to:

- Initialize and finalize the server-side run time
- Implement servants
- Pass parameters and throw exceptions
- Create servants and register them with the Ice run time

Because the mapping for Slice data types is identical for clients and servers, the server-side mapping only adds a few additional mechanisms to the client side: a small API to initialize and finalize the run time, plus a few rules for how to derive servant classes from skeletons and how to register servants with the server-side run time.

Although the examples we present are very simple, they accurately reflect the basics of writing an Ice server. Of course, for more sophisticated servers, you will be using additional APIs, for example, to improve performance or scalability. However, these APIs are all described in Slice, so, to use these APIs, you need not learn any JavaScript mapping rules beyond those we described here.

Support for server-side activities in Ice for JavaScript is currently limited to callbacks over bidirectional connections, which allows a JavaScript client to receive "push notifications" from a server. In this section, we describe the server-related activities of a JavaScript client.

Topics

- Server-Side JavaScript Mapping for Interfaces
- Parameter Passing in JavaScript
- Raising Exceptions in JavaScript
- Object Incarnation in JavaScript
- Asynchronous Method Dispatch (AMD) in JavaScript

Copyright 2016, ZeroC, Inc.

Server-Side JavaScript Mapping for Interfaces

The server-side mapping for interfaces provides an up-call API for the Ice run time: by implementing member functions in a servant class, you provide the hook that gets control flow from the Ice server-side run time into your application code.

On this page:

- Skeleton Types in JavaScript
- Servant Types in JavaScript
  - Normal and idempotent Operations in JavaScript

Skeleton Types in JavaScript

On the client side, interfaces map to proxy types. On the server side, interfaces map to *skeleton* types. A skeleton is a type that conceptually has an abstract member function for each operation on the corresponding interface. For example, consider our Slice definition for the `Node` i nterface:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };
    // ...
};
```

The Slice compiler generates the following definition for this interface:

JavaScript

```
Filesystem.Node = function() { ... };
Filesystem.Node.prototype = new Ice.Object();
Filesystem.Node.prototype.constructor = Filesystem.Node;
// Mapping-internal code here...
```

The important points to note here are:

- As for the client side, Slice modules are mapped to JavaScript scopes with the same name, so the skeleton definitions are part of the `Filesystem` scope.
- For each Slice interface <*`interface-name`*>, the compiler generates a JavaScript type <*`interface-name`*> (`Node` in this example). This type extends `Ice.Object` and serves as the actual skeleton; it is the base type from which you derive your servant implementation.

Servant Types in JavaScript

In order to provide an implementation for an Ice object, you must create a servant type that inherits from the corresponding skeleton. For example, to create a servant for the `Node` interface, you could write:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var NodeI = Ice.Class(Filesystem.Node, {
    __init__: function(name) {
        this._name = name;
    },
    name: function(current) {
        return this._name;
    }
});
```

By convention, servant types have the name of their interface with an `I`-suffix, so the servant type for the `Node` interface is called `NodeI`. (This is a convention only: as far as the Ice run time is concerned, you can choose any name you prefer for your servant types.) Note that `NodeI` extends `Filesystem.Node`, that is, it derives from its skeleton.

Here we're using the `Ice.Class` convenience function to define our servant. This function establishes the necessary prototype relationship with the given base type and adds a method to the new type's prototype for each property in the supplied object. The value of the `__init__` property (if defined) becomes the new type's constructor.

As far as Ice is concerned, the `NodeI` type must implement only a single function: the `name` method defined by the Slice interface. You can add other methods and data members as you see fit to support your implementation. For example, in the preceding definition, we added a `_name` member and a constructor. (Obviously, the constructor initializes the `_name` member and the `name` method returns its value.) You can ignore the `current` parameter for now.

Normal and `idempotent` Operations in JavaScript

Whether an operation is an ordinary operation or an `idempotent` operation has no influence on the way the operation is mapped. To illustrate this, consider the following interface:

Slice

```
interface Example {
               void   normalOp();
    idempotent void   idempotentOp();
    idempotent string readonlyOp();
};
```

The corresponding servant methods look like this:

JavaScript

```
var ExampleI = Ice.Class(Example, {
    normalOp:     function(current) { ... },
    idempotentOp: function(current) { ... },
    readonlyOp:   function(current) { ... }
});
```

Note that the signatures of the member functions are unaffected by the `idempotent` qualifier.

Copyright 2016, ZeroC, Inc.

See Also

- Slice for a Simple File System
- JavaScript Mapping for Interfaces
- Parameter Passing in JavaScript
- Raising Exceptions in JavaScript
- The Current Object

Copyright 2016, ZeroC, Inc.

Parameter Passing in JavaScript

On this page:

- Server-Side Parameters in JavaScript

Server-Side Parameters in JavaScript

For each parameter of a Slice operation, the JavaScript mapping generates a corresponding parameter for the method in the servant. In addition, every operation has an additional, trailing parameter of type `Ice.Current`. For example, the `name` operation of the `Node` interface has no parameters, but the `name` method of the servant has a single parameter of type `Ice.Current`. We will ignore this parameter for now.

For a Slice operation that returns multiple values, the implementation returns them in an array consisting of a non-`void` return value, if any, followed by the `out` parameters in the order of declaration. An operation returning only one value simply returns the value itself.

An operation returns multiple values when it declares multiple out parameters, or when it declares a non-`void` return type and at least one `out` parameter.

To illustrate the rules, consider the following interface that passes string parameters in all possible directions:

Slice

```
module M {
    interface Example {
        string op1(string sin);
        void op2(string sin, out string sout);
        string op3(string sin, out string sout);
    };
};
```

The servant methods would look as follows:

JavaScript

```
var ExampleI = Ice.Class(M.Example, {
    op1: function(sin, current) { ... },
    op2: function(sin, current) { ... },
    op3: function(sin, current) { ... }
});
```

The signatures of the JavaScript methods are identical because they all accept a single `in` parameter, but their implementations differ in the way they return values. For example, we could implement the operations as follows:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var ExampleI = Ice.Class(M.Example, {
    op1: function(sin, current) {
        console.log(sin); // In params are initialized
        return "Done";    // Return value
    },
    op2: function(sin, current) {
        console.log(sin);      // In params are initialized
        return "Hello World!"; // Out parameter
    },
    op3: function(sin, current) {
        console.log(sin); // In params are initialized
        return ["Done", "Hello World!"];
    }
});
```

Notice that `op1` and `op2` return their string values directly, whereas `op3` returns an array consisting of the return value followed by the `out` parameter.

This code is in no way different from what you would normally write if you were to pass strings to and from a function; the fact that remote procedure calls are involved does not impact on your code in any way. The same is true for parameters of other types, such as proxies, classes, or dictionaries: the parameter passing conventions follow normal JavaScript rules and do not require special-purpose API calls.

See Also

- Server-Side JavaScript Mapping for Interfaces
- Raising Exceptions in JavaScript
- The Current Object

Copyright 2016, ZeroC, Inc.

Raising Exceptions in JavaScript

To throw an exception from an operation implementation, you simply instantiate the exception, initialize it, and throw it. For example:

JavaScript

```
// ...
write: function(text, current) {
    try {
        // Try to write file contents here...
    }
    catch(err) {
        var ex = new GenericError("Exception during write operation");
        ex.ice_cause = err; // Not returned to client but may be useful
        throw ex;
    }
}
```

If you throw an arbitrary JavaScript error (such as an instance of `Error`), the Ice run time catches the exception and then returns an `UnknownException` to the client. Similarly, if you throw an "impossible" user exception (a user exception that is not listed in the exception specification of the operation), the client receives an `UnknownUserException`.

If you throw an Ice run-time exception, such as `MemoryLimitException`, the client receives an `UnknownLocalException`. For that reason, you should never throw system exceptions from operation implementations. If you do, all the client will see is an `UnknownLocalException`, which does not tell the client anything useful.

Three run-time exceptions are treated specially and not changed to `UnknownLocalException` when returned to the client: `ObjectNotExistException`, `OperationNotExistException`, and `FacetNotExistException`.

See Also

- Run-Time Exceptions
- JavaScript Mapping for Exceptions
- Server-Side JavaScript Mapping for Interfaces
- Parameter Passing in JavaScript

Copyright 2016, ZeroC, Inc.

Object Incarnation in JavaScript

Having created a servant such as the rudimentary `NodeI` type, you can instantiate the type to create a concrete servant that can receive invocations from a client. However, merely instantiating a servant is insufficient to incarnate an object. Specifically, to provide an implementation of an Ice object, you must take the following steps:

1. Instantiate a servant class.
2. Create an identity for the Ice object incarnated by the servant.
3. Inform the Ice run time of the existence of the servant.
4. Pass a proxy for the object to a client so the client can reach it.

On this page:

- Instantiating a JavaScript Servant
- Creating an Identity in JavaScript
- Activating a JavaScript Servant
- UUIDs as Identities in JavaScript
- Creating Proxies in JavaScript
  - Proxies and Servant Activation in JavaScript
  - Direct Proxy Creation in JavaScript

Instantiating a JavaScript Servant

Instantiating a servant means to allocate an instance:

JavaScript

```
var servant = new NodeI("Fred");
```

This code creates a new `NodeI` instance.

Creating an Identity in JavaScript

Each Ice object requires an identity. That identity must be unique for all servants using the same object adapter.

The Ice object model assumes that all objects (regardless of their adapter) have a globally unique identity.

An Ice object identity is a structure with the following Slice definition:

Slice

```
module Ice {
    struct Identity {
        string name;
        string category;
    };
    // ...
};
```

The full identity of an object is the combination of both the `name` and `category` fields of the `Identity` structure. For now, we will leave the `category` field as the empty string and simply use the `name` field. (The `category` field is most often used in conjunction with servant locators.)

To create an identity, we simply assign a key that identifies the servant to the `name` field of the `Identity` structure:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var id = new Ice.Identity();
id.name = "Fred"; // Not unique, but good enough for now
```

Activating a JavaScript Servant

Merely creating a servant instance does nothing: the Ice run time becomes aware of the existence of a servant only once you explicitly tell the object adapter about the servant. To activate a servant, you invoke the `add` operation on the object adapter. Assuming that we have access to the object adapter in the `_adapter` variable, we can write:

JavaScript

```
this._adapter.add(servant, id);
```

Note the two arguments to `add`: the servant and the object identity. Calling `add` on the object adapter adds the servant and the servant's identity to the adapter's servant map and links the proxy for an Ice object to the correct servant instance in the server's memory as follows:

1. The proxy for an Ice object, apart from addressing information, contains the identity of the Ice object. When a client invokes an operation, the object identity is sent with the request to the server.
2. The object adapter receives the request, retrieves the identity, and uses the identity as an index into the servant map.
3. If a servant with that identity is active, the object adapter retrieves the servant from the servant map and dispatches the incoming request into the correct member function on the servant.

Assuming that the object adapter is in the active state, client requests are dispatched to the servant as soon as you call `add`.

UUIDs as Identities in JavaScript

The Ice object model assumes that object identities are globally unique. One way of ensuring that uniqueness is to use UUIDs (Universally Unique Identifiers) as identities. Ice for JavaScript provides a helper function that we can use to create such identities:

JavaScript

```
var uuid = Ice.generateUUID();
console.log(uuid);
```

When executed, this code prints a unique string such as `5029a22c-e333-4f87-86b1-cd5e0fcce509`. Each call to `generateUUID` creates a string that differs from all previous ones.

You can use a UUID such as this to create object identities. For convenience, the object adapter has an operation `addWithUUID` that generates a UUID and adds a servant to the servant map in a single step. Using this operation, we can create an identity and register a servant with that identity in a single step as follows:

Copyright 2016, ZeroC, Inc.

JavaScript

```
this._adapter.addWithUUID(new NodeI("Fred"));
```

Creating Proxies in JavaScript

Once we have activated a servant for an Ice object, the server can process incoming client requests for that object. However, clients can only access the object once they hold a proxy for the object.

At present, Ice for JavaScript only supports server-side activity over bidirectional connections. There are two ways of using bidirectional connections:

- In conjunction with a Glacier2 router, which requires very little additional code in the client aside from ensuring that the object adapter is correctly configured to use the router. This is the simplest and most common way of employing bidirectional connections.
- Manually configuring a connection for bidirectional use.

The remaining discussion in this section applies when using a Glacier2 router.

Proxies and Servant Activation in JavaScript

The `add` and `addWithUUID` servant activation operations on the object adapter return a proxy for the corresponding Ice object. This means we can write:

JavaScript

```
var proxy = NodePrx.uncheckedCast(this._adapter.addWithUUID(new NodeI("Fred")));
```

Here, `addWithUUID` both activates the servant and returns a proxy for the Ice object incarnated by that servant in a single step.

Note that we need to use an `uncheckedCast` here because `addWithUUID` returns a proxy of type `Ice.ObjectPrx`.

Direct Proxy Creation in JavaScript

The object adapter offers an operation to create a proxy for a given identity:

Slice

```
module Ice {
    local interface ObjectAdapter {
        Object* createProxy(Identity id);
        // ...
    };
};
```

Note that `createProxy` creates a proxy for a given identity whether a servant is activated with that identity or not. In other words, proxies have a life cycle that is quite independent from the life cycle of servants:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var id = new Ice.Identity();
id.name = Ice.generateUUID();
var o = this._adapter.createProxy(id);
```

This creates a proxy for an Ice object with the identity returned by `generateUUID`. Obviously, no servant yet exists for that object so, if we return the proxy to a peer and the peer invokes an operation on the proxy, the peer will receive an `ObjectNotExistException`. (We examine these life cycle issues in more detail in Object Life Cycle.)

See Also

- Hello World Application
- Server-Side JavaScript Mapping for Interfaces
- Object Adapter States
- Servant Locators
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

Asynchronous Method Dispatch (AMD) in JavaScript

JavaScript is single-threaded, therefore it is very important to consider how servant implementation decisions affect the overall application. In a JavaScript program with a graphical user interface, spending too much time in the implementation of a Slice operation can delay the processing of UI events and adversely impact the user experience.

Asynchronous Method Dispatch (AMD), the server-side equivalent of AMI, allows a server to receive a request but then suspend its processing. When processing resumes and the results are available, the server sends a response explicitly using a callback object provided by the Ice run time.

AMD is transparent to the client, that is, there is no way for a client to distinguish a request that, in the server, is processed synchronously from a request that is processed asynchronously.

One use case that requires AMD is nested invocations. Since all outgoing invocations have asynchronous semantics, an operation implementation whose results depend on the outcome of a nested invocation must postpone its response until the nested invocation completes.

An alternate use case for AMD is an operation that requires further processing after completing the client's request. In order to minimize the client's delay, the operation returns the results and then continues to perform additional work.

On this page:

- Enabling AMD with Metadata in JavaScript
- AMD Mapping in JavaScript
  - Callback object for AMD
  - Servant method for AMD
- AMD Exceptions in JavaScript
- AMD Example in JavaScript

Enabling AMD with Metadata in JavaScript

To enable asynchronous dispatch, you must add an `["amd"]` metadata directive to your Slice definitions. The directive applies at the interface and the operation level. If you specify `["amd"]` at the interface level, all operations in that interface use asynchronous dispatch; if you specify `["amd"]` for an individual operation, only that operation uses asynchronous dispatch. In either case, the metadata directive *repl aces* synchronous dispatch, that is, a particular operation implementation must use synchronous or asynchronous dispatch and cannot use both.

Consider the following Slice definitions:

Slice

```
["amd"] interface I {
    bool isValid();
    float computeRate();
};

interface J {
    ["amd"] void startProcess();
    int endProcess();
};
```

In this example, both operations of interface `I` use asynchronous dispatch, whereas, for interface `J`, `startProcess` uses asynchronous dispatch and `endProcess` uses synchronous dispatch.

Specifying metadata at the operation level (rather than at the interface or class level) minimizes complexity: although the asynchronous model is more flexible, it is also more complicated to use. It is therefore in your best interest to limit the use of the asynchronous model to those operations that need it, while using the simpler synchronous model for the rest.

AMD Mapping in JavaScript

The AMD mapping differs from the default synchronous model in two ways:

Copyright 2016, ZeroC, Inc.

1. A callback object is passed as the first argument to a servant method
2. The name of the servant method requires an `_async` prefix

Callback object for AMD

A servant method receives a callback object that it uses to notify the Ice run time about the completion of an operation. The object defines two methods:

JavaScript

```
function ice_response(<params>)
function ice_exception(ex)
```

The `ice_response` function allows the server to report the successful completion of the operation. If the operation's Slice definition has a non-`void` return type, the first parameter to `ice_response` is the return value. Parameters corresponding to the operation's out-parameters follow the return value, in the order of declaration.

The `ice_exception` function allows the server to raise an exception. The Ice run time validates the exception value using the same semantics as for synchronous dispatch.

Neither `ice_response` nor `ice_exception` throws any exceptions to the caller.

Servant method for AMD

The servant method, whose name has the suffix `_async`, always has a `void` return type. The first parameter is the callback object described above. The remaining parameters comprise the in-parameters of the operation, in the order of declaration.

For example, suppose we have defined the following operation:

Slice

```
interface I {
  ["amd"] int foo(short s, out long l);
};
```

The callback object's `ice_response` method has the signature shown below:

JavaScript

```
function ice_response(retval, l)
```

The servant method for asynchronous invocation of operation `foo` would have this signature:

JavaScript

```
foo_async: function(cb, s) { ... };
```


Copyright 2016, ZeroC, Inc.

AMD Exceptions in JavaScript

There are two processing contexts in which the logical implementation of an AMD operation may need to report an exception: the dispatch context (the call sequence that invokes the servant method), and the response context (the call sequence that sends the response).

These are not necessarily two different contexts: it is legal to send the response from the dispatch context.

Although we recommend that the callback object be used to report all exceptions to the client, it is legal for the implementation to raise an exception instead, but only from the dispatch context.

As you would expect, an exception raised from a response context cannot be caught by the Ice run time; the application's run time environment determines how such an exception is handled. Therefore, a response context must ensure that it traps all exceptions and sends the appropriate response using the callback object. Otherwise, if a response context is terminated by an uncaught exception, the request may never be completed and the client might wait indefinitely for a response.

Whether raised in a dispatch context or reported via the callback object, user exceptions are validated and local exceptions may undergo translation.

AMD Example in JavaScript

To demonstrate the use of AMD in Ice, consider an operation that must make a nested invocation:

Slice

```
module Demo {

    exception RequestFailed {
        string reason;
    };

    interface TaxManager {
        float computeTax(float amount)
            throws RequestFailed;
    };

    interface ShoppingCart {
        ["amd"] float computeTotal()
            throws RequestFailed;
    };
};
```

In this over-simplified example, a shopping cart object must query a tax manager object in order to compute the total amount owed by the customer.

Our servant class derives from `Demo.ShoppingCart` and supplies a definition for the `computeTotal_async` method:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var ShoppingCartI = Ice.Class(Demo.ShoppingCart, {
    __init__: function() {
        this._subtotal = 0.0;
    },
    computeTotal_async: function(cb, current) {
        var taxManager = ... // get TaxManager proxy
        var subtotal = this._subtotal;
        taxManager.computeTax(subtotal).then(
            function(tax) {
                cb.ice_response(subtotal + tax);
            },
            function(ex) {
                if(ex instanceof Demo.RequestFailed)
                    cb.ice_exception(ex); // Relay RequestFailed exception from
computeTax
                else
                    cb.ice_exception(new Demo.RequestFailed("failed: " +
ex.toString()));
            });
    }
});
```

The implementation of `computeTotal_async` makes a nested invocation of `computeTax`. The success and failure callbacks use the AMD callback object to complete the invocation of the `computeTotal` operation.

See Also

- User Exceptions
- Run-Time Exceptions
- JavaScript Mapping for Operations

Copyright 2016, ZeroC, Inc.

Objective-C Mapping

Topics

- Client-Side Slice-to-Objective-C Mapping
- Server-Side Slice-to-Objective-C Mapping

Copyright 2016, ZeroC, Inc.

Client-Side Slice-to-Objective-C Mapping

The client-side Slice-to-Objective-C mapping defines how Slice data types are translated to Objective-C types, and how clients invoke operations, pass parameters, and handle errors. Much of the Objective-C mapping is intuitive. For example, Slice dictionaries map to Cocoa framework dictionaries, so there is little new you have to learn in order to use Slice dictionaries in Objective-C.

The Objective-C mapping is thread-safe. For example, you can concurrently invoke operations on an object from different threads without the risk of race conditions or corrupting data structures in the Ice run time, but you must still synchronize access to application data from different threads. For example, if you have two threads sharing a sequence, you cannot safely have one thread insert into the sequence while another thread is iterating over the sequence. However, you only need to concern yourself with concurrent access to your own data — the Ice run time itself is fully thread safe, and none of the Ice API calls require you to acquire or release a lock before you safely can make the call.

Much of what appears in this chapter is reference material. We suggest that you skim the material on the initial reading and refer back to specific sections as needed. However, we recommend that you read at least the mappings for exceptions, interfaces, and operations in detail because these sections cover how to call operations from a client, pass parameters, and handle exceptions.

In order to use the Objective-C mapping, you should need no more than the Slice definition of your application and knowledge of the Objective-C mapping rules. In particular, looking through the generated code in order to discern how to use the Objective-C mapping is likely to be confusing because the generated code is not necessarily meant for human consumption and, occasionally, contains various cryptic constructs to deal with mapping internals. Of course, occasionally, you may want to refer to the generated code to confirm a detail of the mapping, but we recommend that you otherwise use the material presented here to see how to write your client-side code.

ICE, ICEMX, GLACIER2, ICESTORM, ICEGRID Prefixes

All of the APIs for the Ice run time are prefixed by `ICE`, to avoid clashes with definitions for other libraries or applications. Parts of the Ice API are generated from Slice definitions; other parts provide special-purpose definitions that do not have a corresponding Slice definition. Regardless of the way they are defined, the `ICE` prefix universally applies to all entry points in the Ice run time. We will incrementally cover the contents of the Ice API throughout the remainder of the manual.

The APIs for IceMX, Glacier2, IceStorm and IceGrid are also prefixed with respectively ICEMX, GLACIER2, ICESTORM and ICEGRID.

Topics

- Objective-C Mapping for Modules
- Objective-C Mapping for Identifiers
- Objective-C Mapping for Built-In Types
- Objective-C Mapping for Enumerations
- Objective-C Mapping for Structures
- Objective-C Mapping for Sequences
- Objective-C Mapping for Dictionaries
- Objective-C Mapping for Constants
- Objective-C Mapping for Exceptions
- Objective-C Mapping for Interfaces
- Objective-C Mapping for Operations
- Objective-C Mapping for Local Interfaces
- Objective-C Mapping for Classes
- Objective-C Mapping for Interfaces by Value
- Objective-C Mapping for Optional Data Members
- Asynchronous Method Invocation (AMI) in Objective-C
- slice2objc Command-Line Options
- Example of a File System Client in Objective-C

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Modules

Because Objective-C does not support namespaces, a Slice module maps to a prefix for the identifiers defined inside the modules. By default, the prefix is the same as the name of the module:

Slice

```
module example
{
    enum Color { Red, Green, Blue };
};
```

With this definition, the Slice identifier `Color` maps to the Objective-C identifier `exampleColor`.

For nested modules, by default, the module identifiers are concatenated. For example, consider the following Slice definition:

Slice

```
module outer {
    module inner {
        enum Color { Red, Green, Blue };
    };
};
```

With this definition, the Slice identifier `Color` becomes `outerinnerColor` in Objective-C.

You can use a metadata directive to change the default mapping to a different prefix. For example:

Slice

```
["objc:prefix:OUT"]
module outer {
    enum Vehicle { Car, Truck, Bicycle };

    module inner {
        enum Color { Red, Green, Blue };
    };
};
```

With this definition, `Vehicle` maps to `OUTVehicle`. However, `Color` still maps to `outerinnerColor`, that is, the metadata directive applies only to types defined in the `outer` module, but not to types that are defined in nested modules. If you want to assign a prefix for types in the nested module, you must use a separate metadata directive, for example:

Copyright 2016, ZeroC, Inc.

Slice

```
["objc:prefix:OUT"]
module outer {
    enum Vehicle { Car, Truck, Bicycle };

    ["objc:prefix:IN"]
    module inner {
        enum Color { Red, Green, Blue };
    };
};
```

With this definition, `Vehicle` maps to `OUTVehicle`, and `Color` maps to `INColor`.

For the remainder of the examples in this chapter, we assume that Slice definitions are enclosed by a module `Example` that is annotated with the metadata directive `["objc:prefix:EX"]`.

See Also

- Objective-C Mapping for Identifiers
- Objective-C Mapping for Built-In Types
- Objective-C Mapping for Enumerations
- Objective-C Mapping for Structures
- Objective-C Mapping for Sequences
- Objective-C Mapping for Dictionaries
- Objective-C Mapping for Constants
- Objective-C Mapping for Exceptions
- Objective-C Mapping for Interfaces

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Identifiers

Objective-C identifiers are derived from Slice identifiers. The exact Objective-C identifier that is generated depends on the context. For types that are nested in modules (and hence have global visibility in Objective-C), the generated Objective-C identifiers are prefixed with their module prefix. Slice identifiers that do not have global visibility (such as operation names and structure members) do not use the module prefix and are preserved without change. For example, consider the following Slice definition:

Slice

```
["objc:prefix:EX"]
module Example {
    struct Point {
        double x;
        double y;
    };
};
```

This maps to the following Objective-C definition:

Objective-C

```
@interface EXPoint : NSObject <NSCopying>
{
    @private
        ICEDouble x;
        ICEDouble y;
}

@property(nonatomic, assign) ICEDouble x;
@property(nonatomic, assign) ICEDouble y;

// More definitions here...
@end
```

If a Slice identifier is the same as an Objective-C keyword, the corresponding Objective-C identifier has an underscore suffix. For example, Slice `while` maps to Objective-C `while_`.

In some cases, the Objective-C mapping generates more than one identifier for a given Slice construct. For example, an interface `Intf` generates the identifiers `EXIntf` and `EXIntfPrx`. If a Slice identifier happens to be an Objective-C keyword, the underscore suffix applies only where necessary, so an interface `while` generates `EXWhile` and `EXWhilePrx`.

Note that Slice operation and member names can clash with the name of an inherited method, property, or instance variable. For example:

Slice

```
exception Failed {
    string reason; // Clashes with NSException
};
```

This is a legal Slice definition. However, the generated exception class derives from `NSException`, which defines a `reason` method. To

Copyright 2016, ZeroC, Inc.

avoid hiding the method in the base class, the generated exception class maps the Slice `reason` member to the Objective-C property `reason_`, just as it would for a keyword.

This escape mechanism applies to all generated classes that directly or indirectly derive from `NSObject` or `NSException`.

Internal Identifiers in Objective-C

Any methods that contain two or more adjacent underscores (such as `read__` and `op____`) are internal to the Objective-C mapping implementation and are not for use by application code.

See Also

- Objective-C Mapping for Modules
- Objective-C Mapping for Built-In Types
- Objective-C Mapping for Enumerations
- Objective-C Mapping for Structures
- Objective-C Mapping for Sequences
- Objective-C Mapping for Dictionaries
- Objective-C Mapping for Constants
- Objective-C Mapping for Exceptions
- Objective-C Mapping for Interfaces

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Built-In Types

The Slice built-in types are mapped to Objective-C types as shown below.

Slice	Objective-C
`bool`	`BOOL`
`byte`	`ICEByte`
`short`	`ICEShort`
`int`	`ICEInt`
`long`	`ICELong`
`float`	`ICEFloat`
`double`	`ICEDouble`
`string`	`NSString` or `NSMutableString`

Slice `bool` maps to Objective-C `BOOL`. The remaining integral and floating-point types map to Objective-C type definitions instead of native types. This allows the Ice run time to provide a definition as appropriate for each target architecture. (For example, `ICELong` might be defined as `long` on one architecture and as `long long` on another.)

Note that `ICEByte` is a typedef for `unsigned char`. This guarantees that byte values are always in the range 0..255, and it ensures that right-shifting an `ICEByte` does not cause sign-extension.

Whether a Slice string maps to `NSString` or `NSMutableString` depends on the context. `NSMutableString` is used in some cases for operation parameters; otherwise, if a string is a data member of a Slice structure, class, or exception, it maps to `NSString`. (We will discuss these differences in more detail as we cover the mapping of the relevant Slice language features.)
See Also

- Objective-C Mapping for Modules
- Objective-C Mapping for Identifiers
- Objective-C Mapping for Enumerations
- Objective-C Mapping for Structures
- Objective-C Mapping for Sequences
- Objective-C Mapping for Dictionaries
- Objective-C Mapping for Constants
- Objective-C Mapping for Exceptions
- Objective-C Mapping for Interfaces

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Enumerations

A Slice enumeration maps to the corresponding enumeration in Objective-C. For example:

Slice

```
["objc:prefix:EX"]
module Example {
    enum Fruit { Apple, Pear, Orange };
};
```

The generated Objective-C definition is:

Objective-C

```
typedef enum {
    EXApple, EXPear, EXOrange
} EXFruit;
```

See Also

- Objective-C Mapping for Modules
- Objective-C Mapping for Identifiers
- Objective-C Mapping for Built-In Types
- Objective-C Mapping for Structures
- Objective-C Mapping for Sequences
- Objective-C Mapping for Dictionaries
- Objective-C Mapping for Constants
- Objective-C Mapping for Exceptions
- Objective-C Mapping for Interfaces

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Structures

On this page:

- Basic Objective-C Mapping for Structures
- Mapping for Data Members in Objective-C
- Creating and Initializing Structures in Objective-C
- Copying Structures in Objective-C
- Deallocating Structures in Objective-C
- Structure Comparison and Hashing in Objective-C

Basic Objective-C Mapping for Structures

A Slice structure maps to an Objective-C class.

For each Slice data member, the generated Objective-C class has a corresponding property. For example, here is our `Employee` structure once more:

Slice

```
struct Employee {
    long number;
    string firstName;
    string lastName;
};
```

The Slice-to-Objective-C compiler generates the following definition for this structure:

Objective-C

```
@interface EXEmployee : NSObject <NSCopying>
{
    @private
        ICELong number;
        NSString *firstName;
        NSString *lastName;
}

@property(nonatomic, assign) ICELong number;
@property(nonatomic, retain) NSString *firstName;
@property(nonatomic, retain) NSString *lastName;

-(id) init:(ICELong)number firstName:(NSString *)firstName
                          lastName:(NSString *)lastName;
+(id) employee:(ICELong)number firstName:(NSString *)firstName
                              lastName:(NSString *)lastName;
+(id) employee;
// This class also overrides copyWithZone,
// hash, isequal, and dealloc.
@end
```

Mapping for Data Members in Objective-C

Copyright 2016, ZeroC, Inc.

For each data member in the Slice definition, the Objective-C class contains a corresponding private instance variable of the same name, as well as a property definition that allows you to set and get the value of the corresponding instance variable. For example, given an instance of `EXEmployee`, you can write the following:

Objective-C

```
ICELong number;
EXemployee *e = ...;
[e setNumber:99];
number = [e number];

// Or, more concisely with dot notation:

e.number = 99;
number = e.number;
```

Properties that represent data members always use the `nonatomic` property attribute. This avoids the overhead of locking each data member during access. The second property attribute is `assign` for integral and floating-point types and `retain` for all other types (such as strings, structures, and so on.)

Note that, for types that have immutable and mutable variants (strings, sequences, and dictionaries), the corresponding data member uses the immutable variant. This allows the application to assign an immutable object to the data member. You can safely cast the data member to the mutable variant if the structure was created by the Ice run time: the unmarshaling code always creates and assigns the mutable version to the data member.

Creating and Initializing Structures in Objective-C

Structures provide the typical (inherited) `init` method:

Objective-C

```
EXEmployee *e = [[EXEmployee alloc] init];
// ...
[e release];
```

As usual, `init` initializes the instance variables of the structure with zero-filled memory, with the following exceptions:

- A string data member is initialized to an empty string
- An enumerator data member is initialized to the first enumerator
- A structure data member is initialized to a default-constructed value

You can also declare default values in your Slice definition, in which case this `init` method initializes each data member with its declared value instead.

In addition, a structure provides a second `init` method that accepts one parameter for each data member of the structure:

Objective-C

```
-(id) init:(ICELong)number firstName:(NSString *)firstName
                          lastName:(NSString *)lastName;
```

Note that the first parameter is always unlabeled; the second and subsequent parameters have a label that is the same as the name of the

Copyright 2016, ZeroC, Inc.

corresponding Slice data member. The additional `init` method allows you to instantiate a structure and initialize its data members in a single statement:

Objective-C

```
EXEmployee *e = [[EXEmployee alloc] init:99 firstName:@"Brad" lastName:@"Cox"];
// ...
[e release];
```

`init` applies the memory management policy of the corresponding properties, that is, it calls `retain` on the `firstName` and `lastName` ar guments.

Each structure also provides two convenience constructors that mirror the `init` methods: a parameter-less convenience constructor and one that has a parameter for each Slice data member:

Objective-C

```
+(id) employee;
+(id) employee:(ICELong)number firstName:(NSString *)firstName
                               lastName:(NSString *)lastName;
```

The convenience constructors have the same name as the mapped Slice structure (without the module prefix). As usual, they allocate an instance, perform the same initialization actions as the corresponding `init` methods, and call `autorelease` on the return value:

Objective-C

```
EXEmployee *e = [EXEmployee employee:99 firstName:@"Brad" lastName:@"Cox"];

// No need to call [e release] here.
```

Copying Structures in Objective-C

Structures implement the `NSCopying` protocol. Structures are copied by assigning instance variables of value type and calling `retain` on each instance variable of non-value type. In other words, the copy is shallow:

Objective-C

```
EXEmployee *e = [EXEmployee employee:99 firstName:@"Brad" lastName:@"Cox"];
EXEmployee *e2 = [e copy];
NSAssert(e.number == e2.number);
NSAssert([e.firstName == e2.firstName]); // Same instance
// ...
[e2 release];
```

Note that, if you assign an `NSMutableString` to a structure member and use the structure as a dictionary key, you must not modify the string inside the structure without copying it because doing so will corrupt the dictionary.

Copyright 2016, ZeroC, Inc.

Deallocating Structures in Objective-C

Each structure implements a `dealloc` method that calls `release` on each instance variable with a `retain` property attribute. This means that structures take care of the memory management of their contents: releasing a structure automatically releases all its instance variables.

Structure Comparison and Hashing in Objective-C

Structures implement `isEqual`, so you can compare them for equality. Two structures are equal if all their instance variables are equal. For value types, equality is determined by the == operator; for non-value types other than classes, equality is determined by the corresponding instance variable's `isEqual` method. Classes are compared by comparing their identity: two class members are equal if they both point at the same instance.

The `hash` method returns a hash value is that is computed from the hash value of all of the structure's instance variables.

See Also

- Structures
- Dictionaries
- Objective-C Mapping for Modules
- Objective-C Mapping for Identifiers
- Objective-C Mapping for Built-In Types
- Objective-C Mapping for Enumerations
- Objective-C Mapping for Sequences
- Objective-C Mapping for Dictionaries
- Objective-C Mapping for Constants
- Objective-C Mapping for Exceptions
- Objective-C Mapping for Interfaces
- Objective-C Mapping for Classes

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Sequences

The Objective-C mapping uses different mappings for sequences of value types (such as `sequence<byte>`) and non-value types (such as `sequence<string>`).

On this page:

- Mapping for Sequences of Value Types in Objective-C
- Mapping of Sequences of Non-Value Types in Objective-C

Mapping for Sequences of Value Types in Objective-C

The following Slice types are value types:

- Integral types (`bool`, `byte`, `short`, `int`, `long`)
- Floating point types (`float`, `double`)
- Enumerated types

Sequences of these types map to a type definition. For example:

Slice

```
enum Fruit { Apple, Pear, Orange };

sequence<byte> ByteSeq;
sequence<int> IntSeq;
sequence<Fruit> FruitSeq;
```

The three Slice sequences produce the following Objective-C definitions:

Objective-C

```
typedef enum {
    EXApple, EXPear, EXOrange
} EXFruit;

typedef NSData EXByteSeq;
typedef NSMutableData EXMutableByteSeq;

typedef NSData EXIntSeq;
typedef NSMutableData EXMutableIntSeq;

typedef NSData EXFruitSeq;
typedef NSMutableData EXMutableFruitSeq;
```

As you can see, each sequence definition creates a pair of type definitions, an immutable version named <*`module-prefix`*><*`Slice-name`*>, and a mutable version named <*`module-prefix`*>`Mutable`<*`Slice-name`*>. This constitutes the entire public API for sequences of value types, that is, sequences of value types simply map to `NSData` or `NSMutableData`. The `NS(Mutable)Data` sequences contain an array of the corresponding element type in their internal byte array.

We chose to map sequences of value types to `NSData` instead of `NSArray` because of the large overhead of placing each sequence element into an `NSNumber` container instance.

For example, here is how you could initialize a byte sequence of 1024 elements with values that are the modulo 128 of the element index in reverse order:

Copyright 2016, ZeroC, Inc.

Objective-C

```
int limit = 1024;
EXMutableByteSeq *bs = [NSMutableData dataWithLength:limit];
ICEByte *p = (ICEByte *)[bs bytes];
while (--limit > 0) {
    *p++ = limit % 0x80;
}
```

Naturally, you do not need to initialize the sequence using a loop. For example, if the data is available in a buffer, you could use the `dataWithBytes:length` or `dataWithBytesNoCopy:length` methods of `NSData` instead.

Here is one way to retrieve the bytes of the sequence:

Objective-C

```
const ICEByte* p = (const ICEByte *)[bs bytes];
const ICEByte* limitp = p + [bs length];
while (p < limitp) {
    printf("%d\n", *p++);
}
```

For sequences of types other than `byte` or `bool`, you must keep in mind that the length of the `NSData` array is not the same as the number of elements. The following example initializes an integer sequence with the first few primes and prints out the contents of the sequence:

Objective-C

```
const int primes[] = { 1, 2, 3, 5, 7, 9, 11, 13, 17, 19, 23 };
EXMutableIntSeq *is = [NSMutableData dataWithBytes:primes length:sizeof(primes)];

const ICEInt *p = (const ICEInt *)[is bytes];
int limit = [is length] / sizeof(*p);
int i;
for(i = 0; i < limit; ++i) {
    printf("%d\n", p[i]);
}
```

The code to manipulate a sequence of enumerators is very similar. For portability, you should not assume a particular size for enumerators. That is, instead of relying on all enumerators having the size of, for example, an `int`, it is better to use `sizeof(EXFruit)` to ensure that you are not overstepping the bounds of the sequence.

Mapping of Sequences of Non-Value Types in Objective-C

Sequences of non-value types, such as sequences of `string`, structures, classes, and so on, map to mutable and immutable type definitions of `NSArray`. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
sequence<string> Page;
sequence<Page> Book;
```

This maps to:

Objective-C

```
typedef NSArray EXPage;
typedef NSMutableArray EXMutablePage;

typedef NSArray EXBook;
typedef NSMutableArray EXMutableBook;
```

You use such sequences as you would use any other `NSArray` in your code. For example:

Objective-C

```
EXMutablePage *page1 = [NSArray arrayWithObjects:
                                @"First line of page one",
                                @"Second line of page one",
                                nil];

EXMutablePage *page2 = [NSArray arrayWithObjects:
                                @"First line of page two",
                                @"Second line of page two",
                                nil];

EXMutableBook *book = [NSMutableArray array];
[book addObject:page1];
[book addObject:page2];
[book addObject:[NSArray array]]; // Empty page
```

This creates a book with three pages; the first two pages contain two lines each, and the third page is empty. You can print the contents of the book as follows:

Copyright 2016, ZeroC, Inc.

Objective-C

```
int pageNum = 0;
for (EXPage *page in book) {
    ++pageNum;
    int lineNum = 0;
    if ([page count] == 0) {
        printf("page %d: <empty>\n", pageNum);
    } else {
        for (NSString *line in page) {
            ++lineNum;
            printf("page %d, line %d: %s\n", pageNum, lineNum, [line UTF8String]);
        }
    }
}
```

This prints:

```
page 1, line 1: First line of page one
page 1, line 2: Second line of page one
page 2, line 1: First line of page two
page 2, line 2: Second line of page two
page 3: <empty>
```

If you have a sequence of proxies or a sequence of classes, to transmit a null proxy or class inside a sequence, you must insert an `NSNull` value into the `NSArray`. In addition, the mapping also allows you to use `NSNull` as the element value of an `NSArray` for elements of type string, structure, sequence, or dictionary. For example, instead of inserting an empty `NSArray` into the book sequence in the preceding example, we could also have inserted `NSNull`:

Objective-C

```
EXMutableBook *book = [NSMutableArray array];
[book addObject:page1];
[book addObject:page2];
[book addObject:[NSNull null]]; // Empty page
```

See Also

- Objective-C Mapping for Modules
- Objective-C Mapping for Identifiers
- Objective-C Mapping for Built-In Types
- Objective-C Mapping for Enumerations
- Objective-C Mapping for Structures
- Objective-C Mapping for Dictionaries
- Objective-C Mapping for Constants
- Objective-C Mapping for Exceptions
- Objective-C Mapping for Interfaces

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Dictionaries

Here is the definition of our `EmployeeMap` once more:

Slice

```
dictionary<long, Employee> EmployeeMap;
```

The following code is generated for this definition:

Objective-C

```
typedef NSDictionary EXEmployeeMap;
typedef NSMutableDictionary EXMutableEmployeeMap;
```

Similar to sequences, Slice dictionaries map to type definitions for `NSDictionary` and `NSMutableDictionary`, with the names <*module-prefix*><*Slice-name*> and <*module-prefix*>Mutable<*Slice-name*>.

As a result, you can use the dictionary like any other `NSDictionary`, for example:

Objective-C

```
EXMutableEmployeeMap *em = [EXMutableEmployeeMap dictionary];
EXEmployee *e = [EXEmployee employee];
e.number = 42;
e.firstName = @"Stan";
e.lastName = @"Lippman";
[em setObject:e forKey:[NSNumber numberWithLong:e.number]];

e = [EXEmployee employee];
e.number = 77;
e.firstName = @"Herb";
e.lastName = @"Sutter";
[em setObject:e forKey:[NSNumber numberWithLong:e.number]];
```

To put a value type into a dictionary (either as the key or the value), you must use `NSNumber` as the object to hold the value. If you have a dictionary that uses a Slice enumeration as the key or the value, you must insert the enumerator as an `NSNumber` that holds an `int`.

To insert a null proxy or null class instance into a dictionary as a value, you must insert `NSNull`.

As a convenience feature, the Objective-C mapping also allows you to insert `NSNull` as the value of a dictionary if the value type of the dictionary is a string, structure, sequence, or dictionary. If you send such a dictionary to a receiver, the Ice run time marshals an empty string, default-initialized structure, empty sequence, or empty dictionary as the corresponding value to the receiver, respectively.
See Also

- Dictionaries
- Objective-C Mapping for Modules
- Objective-C Mapping for Identifiers
- Objective-C Mapping for Built-In Types
- Objective-C Mapping for Enumerations
- Objective-C Mapping for Structures

Copyright 2016, ZeroC, Inc.

- Objective-C Mapping for Sequences
- Objective-C Mapping for Constants
- Objective-C Mapping for Exceptions
- Objective-C Mapping for Interfaces

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Constants

Slice constant definitions map to corresponding Objective?C constant definitions. For example:

Slice

```
const bool      AppendByDefault = true;
const byte      LowerNibble = 0x0f;
const string    Advice = "Don't Panic!";
const short     TheAnswer = 42;
const double    PI = 3.1416;

enum Fruit { Apple, Pear, Orange };
const Fruit     FavoriteFruit = Pear;
```

Here are the generated definitions for these constants:

Objective-C

```
static const BOOL EXAppendByDefault = YES;
static const ICEByte EXLowerNibble = 15;
static NSString * const EXAdvice = @"Don't Panic!";
static const ICEShort EXTheAnswer = 42;
static const ICEDouble EXPI = 3.1416;

typedef enum {
    EXApple, EXPear, EXOrange
} EXFruit;
static const EXFruit EXFavoriteFruit = EXPear;
```

All constants are initialized directly in the generated header file, so they are compile-time constants and can be used in contexts where a compile-time constant expression is required, such as to dimension an array or as the `case` label of a `switch` statement.

Slice string literals that contain non-ASCII characters or universal character names are mapped to Objective-C string literals with these characters replaced by their UTF-8 encoding as octal escapes. For example:

Slice

```
const string Egg = "œuf";
const string Heart = "c\u0153ur";
const string Banana = "\U0001F34C";
```

is mapped to:

Objective-C

```
static NSString * const demoEgg = @"\305\223uf";
static NSString * const demoHeart = @"c\305\223ur";
static NSString * const demoBanana = @"\360\237\215\214";
```


Copyright 2016, ZeroC, Inc.

See Also

- Constants and Literals
- Objective-C Mapping for Modules
- Objective-C Mapping for Identifiers
- Objective-C Mapping for Built-In Types
- Objective-C Mapping for Enumerations
- Objective-C Mapping for Structures
- Objective-C Mapping for Sequences
- Objective-C Mapping for Dictionaries
- Objective-C Mapping for Exceptions
- Objective-C Mapping for Interfaces

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Exceptions

This page describes the Objective-C mapping for exceptions.

On this page:

- Exception Inheritance Hierarchy in Objective-C
- Mapping for Exception Data Members in Objective-C
- Objective-C Mapping for User Exceptions
- Objective-C Mapping for Run-Time Exceptions
    - Creating and Initializing Run-Time Exceptions in Objective-C
- Copying and Deallocating Exceptions in Objective-C
- Exception Comparison and Hashing in Objective-C

Exception Inheritance Hierarchy in Objective-C

Here again is a fragment of the Slice definition for our world time server:

Slice

```
exception GenericError {
    string reason;
};
exception BadTimeVal extends GenericError {};
exception BadZoneName extends GenericError {};
```

These exception definitions map as follows:

Objective-C

```
@interface EXGenericError : ICEUserException
{
@private
    NSString *reason_;
}

@property(nonatomic, retain) NSString *reason_;

// ...
@end

@interface EXBadTimeVal : EXGenericError
// ...
@end

@interface EXBadZoneName : EXGenericError
// ...
@end
```

Each Slice exception is mapped to an Objective-C class. For each exception member, the corresponding class contains a private instance variable and a property. (Obviously, because `BadTimeVal` and `BadZoneName` do not have members, the generated classes for these exceptions also do not have members.)

The inheritance structure of the Slice exceptions is preserved for the generated classes, so `EXBadTimeVal` and `EXBadZoneName` inherit

Copyright 2016, ZeroC, Inc.

from `EXGenericError`.

In turn, `EXGenericError` derives from `ICEUserException`:

Objective-C

```
@interface ICEException : NSException
-(NSString* *)ice_name;
@end

@interface ICEUserException : ICEException
// ...
@end

@interface ICELocalException : ICEException
// ...
@end
```

Note that `ICEUserException` itself derives from `ICEException`, which derives from `NSException`. Similarly, run-time exceptions derive from a common base class `ICELocalException` that derives from `ICEException`, so we have the inheritance structure shown below:

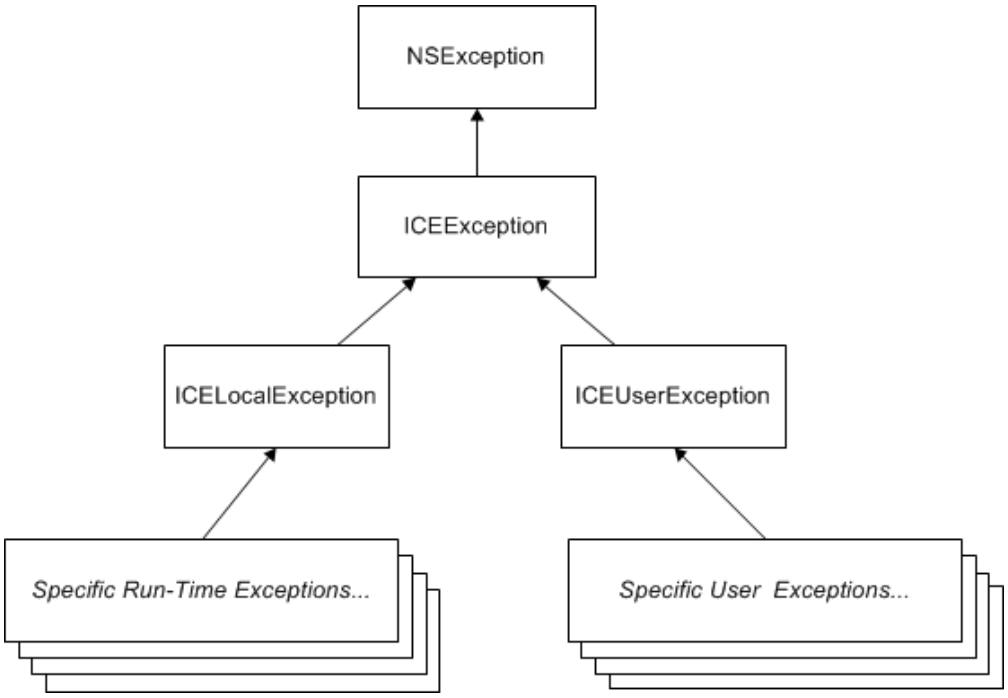

Inheritance structure for exceptions.

`ICEException` provides a single method, `ice_name`, that returns the Slice type ID of the exception with the leading `::` omitted. For example, the return value of `ice_name` for our Slice `GenericError` is `Example::GenericError`.

Mapping for Exception Data Members in Objective-C

As we mentioned earlier, each data member of a Slice exception generates a corresponding Objective-C property. Here is an example that extends our `GenericError` with yet another exception:

Copyright 2016, ZeroC, Inc.

Slice

```
exception GenericError {
    string reason;
};

exception FileError extends GenericError {
    string name;
    int errorCode;
};
```

The generated properties for these exceptions are as follows:

Objective-C

```
@interface EXGenericError : ICEUserException
{
@private
    NSString *reason_;
}

@property(nonatomic, retain) NSString *reason_;

// ...
@end

@interface EXFileError : EXGenericError
{
@private
    NSString *name_;
    ICEInt errorCode;
}

@property(nonatomic, retain) NSString *name_;
@property(nonatomic, assign) ICEInt errorCode;

// ...
@end
```

This is exactly the same mapping as for structure members, with one difference: the `name` and `reason` members map to `name_` and `reason_` properties, whereas — as for structures — `errorCode` maps to `errorCode`. The trailing underscore for `reason_` and `name_` prevents a name collision with the `name` and `reason` methods that are defined by `NSException`: if you call the `name` method, you receive the name that is stored by `NSException`; if you call the `name_` method, you receive the value of the `name_` instance variable of `EXFileError`:

Copyright 2016, ZeroC, Inc.

Objective-C

```
@try {
    // Do something that can throw ExFileError...
}
@catch(EXFileError *ex)
{
    // Print the value of the Slice reason, name,
    // and errorCode members.
    printf("reason: %s, name: %s, errorCode: %d\n",
            [ex.reason_ UTF8String],
            [ex.name_ UTF8String],
            ex.errorCode);

    // Print the NSException name.
    printf("NSException name: %s\n", [[ex name] UTF8String]);
}
```

The same escape mechanism applies if you define exception data members named `callStackReturnAddresses`, `raise`, or `userInfo`.

Objective-C Mapping for User Exceptions

Initialization of exceptions follows the same pattern as for structures: each exception (apart from the inherited no-argument `init` method) provides an `init` method that accepts one argument for each data member of the exception, and two convenience constructors. For example, the generated methods for our `EXGenericError` exception look as follows:

Objective-C

```
@interface EXGenericError : ICEUserException
// ...

-(id) init:(NSString *)reason;
+(id) genericError;
+(id) genericError:(NSString *)reason;
@endif
```

If a user exception has no data members (and its base exceptions do not have data members either), only the inherited `init` method and the no-argument convenience constructor are generated.

The no-argument `init` method and the no-argument convenience constructor initialize the instance variables of the exception with zero-filled memory, with the following exceptions:

- A string data member is initialized to an empty string
- An enumerator data member is initialized to the first enumerator
- A structure data member is initialized to a default-constructed value

If you declare default values in your Slice definition, the inherited `init` method and the no-argument convenience constructor initialize each data member with its declared value instead.

If an exception has a base exception with data members, its `init` method and convenience constructor accept one argument for each Slice data member, in base-to-derived order. For example, here are the methods for the `FileError` exception we defined above:

Copyright 2016, ZeroC, Inc.

Objective-C

```
@interface EXFileError : EXGenericError
// ...

-(id) init:(NSString *)reason name_:(NSString *)name
                           errorCode:(ICEInt)errorCode;
+(id) fileError;
+(id) fileError:(NSString *)reason name_:(NSString *)name
                                errorCode:(ICEInt)errorCode;
@end
```

Note that `init` and the second convenience constructor accept three arguments; the first initializes the `EXGenericError` base, and the remaining two initialize the instance variables of `EXFileError`.

Objective-C Mapping for Run-Time Exceptions

The Ice run time throws run-time exceptions for a number of pre-defined error conditions. All run-time exceptions directly or indirectly derive from `ICELocalException` which, in turn, derives from `ICEException`. (See the above illustration for an example of an inheritance diagram.)

By catching exceptions at the appropriate point in the hierarchy, you can handle exceptions according to the category of error they indicate:

- `NSException`
  This is the root of the complete inheritance tree. Catching `NSException` catches all exceptions, whether they relate to Ice or the Cocoa framework.
- `ICEException`
  Catching `ICEException` catches both user and run-time exceptions.
- `ICEUserException`
  This is the root exception for all user exceptions. Catching `ICEUserException` catches all user exceptions (but not run-time exceptions).
- `ICELocalException`
  This is the root exception for all run-time exceptions. Catching `ICELocalException` catches all run-time exceptions (but not user exceptions).
- `ICETimeoutException`
  This is the base exception for both operation-invocation and connection-establishment timeouts.
- `ICEConnectTimeoutException`
  This exception is raised when the initial attempt to establish a connection to a server times out.

For example, an `ICEConnectTimeoutException` can be handled as `ICEConnectTimeoutException`, `ICETimeoutException`, `ICELocalException`, `ICEException`, or `NSException`.

You will probably have little need to catch run-time exceptions as their most-derived type and instead catch them as `ICELocalException`; the fine-grained error handling offered by the remainder of the hierarchy is of interest mainly in the implementation of the Ice run time. Exceptions to this rule are the exceptions related to facet and object life cycles, which you may want to catch explicitly. These exceptions are `ICEFacetNotExistException` and `ICEObjectNotExistException`, respectively.

Creating and Initializing Run-Time Exceptions in Objective-C

`ICELocalException` provides two properties that return the file name and line number at which an exception was raised:

Copyright 2016, ZeroC, Inc.

Objective-C

```
@interface ICELocalException : ICEException
{
// ...

@property(nonatomic, readonly) NSString* file;
@property(nonatomic, readonly) int line;

-(id)init:(const char*)file line:(int)line;
+(id)localException:(const char*)file line:(int)line;
@end
```

The `init` method and the convenience constructor accept the file name and line number as arguments.

Concrete run-time exceptions that derived from `ICEException` provide a corresponding `init` method and convenience constructor. For example, here is the Slice definition of `ObjectNotExistException`:

Slice

```
local exception RequestFailedException {
    Identity id;
    string facet;
    string operation;
};

local exception ObjectNotExistException extends RequestFailedException {};
```

The Objective-C mapping for `ObjectNotExistException` is:

Copyright 2016, ZeroC, Inc.

Objective-C

```
@interface ICEObjectNotExistException : ICERequestFailedException
// ...
-(id) init:(const char*)file__p line:(int)line__p;
-(id) init:(const char*)file__p
          line:(int)line__p
          id_:(ICEIdentity *)id_
          facet:(NSString *)facet
          operation:(NSString *)operation;
+(id) objectNotExistException:(const char*)file__p
                              line:(int)line__p;
+(id) objectNotExistException:(const char*)file__p
                              line:(int)line__p
                              id_:(ICEIdentity *)id_
                              facet:(NSString *)facet
                              operation:(NSString *)operation;
@end
```

In other words, as for user exceptions, run-time exceptions provide `init` methods and convenience constructors that accept arguments in base-to-derived order. This means that all run-time exceptions require a file name and line number when they are instantiated. For example, you can throw an `ICEObjectNotExistException` as follows:

Objective-C

```
@throw [ICEObjectNotExistException objectNotExistException:__FILE__ line:__LINE__];
```

If you throw this exception in the context of an executing operation on the server side, the `id_`, `facet`, and `operation` instance variables are automatically initialized by the Ice run time.

When you instantiate a run-time exception, the base `NSException` is initialized such that its `name` method returns the same string as `ice_name`; the `reason` and `userInfo` methods return `nil`.

Copying and Deallocating Exceptions in Objective-C

User exceptions and run-time exceptions implement the `NSCopying` protocol, so you can copy them. The semantics are the same as for structures.

Similarly, like structures, exceptions implement a `dealloc` method that takes care of deallocating the instance variables when an exception is released.

Exception Comparison and Hashing in Objective-C

Exceptions do not override `isEqual` or `hash`, so these methods have the behavior inherited from `NSObject`.

See Also

- User Exceptions
- Run-Time Exceptions
- Objective-C Mapping for Modules
- Objective-C Mapping for Identifiers
- Objective-C Mapping for Built-In Types
- Objective-C Mapping for Enumerations
- Objective-C Mapping for Structures

Copyright 2016, ZeroC, Inc.

- Objective-C Mapping for Sequences
- Objective-C Mapping for Dictionaries
- Objective-C Mapping for Constants
- Objective-C Mapping for Interfaces
- Versioning
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Interfaces

The mapping of Slice interfaces revolves around the idea that, to invoke a remote operation, you call a member function on a local class instance that represents the remote object. This makes the mapping easy and intuitive to use because, for all intents and purposes (apart from error semantics), making a remote procedure call is no different from making a local procedure call.

On this page:

- Proxy Classes and Proxy Protocols in Objective-C
- Proxy Instantiation and Casting in Objective-C
  - Using a Checked Cast in Objective-C
  - Using an Unchecked Cast in Objective-C
- Using Proxy Methods in Objective-C
- Object Identity and Proxy Comparison in Objective-C

Proxy Classes and Proxy Protocols in Objective-C

On the client side, interfaces map to a protocol with member functions that correspond to the operations on those interfaces. Consider the following simple interface:

Slice

```
["objc:prefix:EX"]
module Example {
    interface Simple {
        void op();
    }
};
```

The Slice compiler generates the following definitions for use by the client:

Objective-C

```
@interface EXSimplePrx : ICEObjectPrx
// Mapping-internal methods here...
@end

@protocol EXSimplePrx <ICEObjectPrx>
-(void) op;
-(void) op:(ICEContext *)context;
@end;
```

As you can see, the compiler generates a proxy protocol `EXSimplePrx` and a proxy class `EXSimplePrx`. In general, the generated name for both protocol and class is <*`module-prefix`*><*`interface-name`*>`Prx`.

In the client's address space, an instance of `EXSimplePrx` is the local ambassador for a remote instance of the `Simple` interface in a server and is known as a *proxy class instance*. All the details about the server-side object, such as its address, what protocol to use, and its object identity are encapsulated in that instance.

Note that `EXSimplePrx` derives from `ICEObjectPrx`, and that `EXSimplePrx` adopts the `ICEObjectPrx` protocol. This reflects the fact that all Slice interfaces implicitly derive from `Ice::Object`. For each operation in the interface, the proxy protocol has two methods whose name is derived from the operation. For the preceding example, we find that the operation `op` is mapped to two methods, `op` and `op:`.

The second method has a trailing parameter of type `ICEContext`. This parameter is for use by the Ice run time to store information about how to deliver a request; normally, you do not need to supply a value here and can pretend that the trailing parameter does not exist. (We examine the `ICEContext` parameter in detail in Request Contexts. The parameter is also used by IceStorm.)

Proxy Instantiation and Casting in Objective-C

Copyright 2016, ZeroC, Inc.

Client-side application code never manipulates proxy class instances directly. In fact, you are not allowed to instantiate a proxy class directly. Instead, proxy instances are always instantiated on behalf of the client by the Ice run time, so client code never has any need to instantiate a proxy directly.

Proxies are immutable: once the run time has instantiated a proxy, that proxy continues to denote the same remote object and cannot be changed. This means that, if you want to keep a copy of a proxy, it is sufficient to call `retain` on the proxy. (You can also call `copy` on a proxy because `ICEObjectPrx` implements `NSCopying`. However, calling `copy` has the same effect as calling `retain`.)

Proxies are always passed and returned as type `id<<module-prefix><interface-name>Prx>`. For example, for the preceding `Simple` interface, the proxy type is `id<EXSimplePrx>`.

The `ICEObjectPrx` base class provides class methods that allow you to cast a proxy from one type to another, as described below.

Using a Checked Cast in Objective-C

A `checkedCast` tests whether the object denoted by a proxy implements the specified interface:

Objective-C

```
+(id) checkedCast:(id<ICEObjectPrx>)proxy;
```

If so, the cast returns a proxy to the specified interface; otherwise, if the object denoted by the proxy does not implement the specified interface, the cast returns `nil`. Checked casts are typically used to safely down-cast a proxy to a more derived interface. For example, assuming we have Slice interfaces `Base` and `Derived`, you can write the following:

Objective-C

```
id<EXBasePrx> base = ...;  // Initialize base proxy
id<EXDerivedPrx> derived = [EXDerivedPrx checkedCast:base];
if(derived != nil)
{
    // base implements run-time type Derived
    // use derived...
} else {
    // Base has some other, unrelated type
}
```

The expression `[EXDerivedPrx checkedCast:base]` tests whether `base` points at an object of type `Derived` (or an object with a type that is derived from `Derived`). If so, the cast succeeds and `derived` is set to point at the same object as `base`. Otherwise, the cast fails and `derived` is set to `nil`. (Proxies that "point nowhere" are represented by `nil`.)

Calling `checkedCast` on a proxy that is already of the desired proxy type returns immediately that proxy. Otherwise, `checkedCast` always calls `ice_isA` on the target object, and upon success, creates a new instance of the desired proxy class.

The message effectively asks the server "is the object denoted by this proxy of type `Derived`?" The reply from the server is communicated to the application code in form of a successful (non-`nil`) or unsuccessful (`nil`) result. Sending a remote message is necessary because, as a rule, there is no way for the client to find out what the actual run-time type of a proxy is without confirmation from the server. (For example, the server may replace the implementation of the object for an existing proxy with a more derived one.) This means that you have to be prepared for a `checkedCast` to fail. For example, if the server is not running, you will receive an `ICEConnectionRefusedException`; if the server is running, but the object denoted by the proxy no longer exists, you will receive an `ICEObjectNotExistException`.

Using an Unchecked Cast in Objective-C

In some cases, it is known that an object supports a more derived interface than the static type of its proxy. For such cases, you can use an unchecked down-cast:

Copyright 2016, ZeroC, Inc.

Objective-C

```
+(id) uncheckedCast:(id<ICEObjectPrx>)proxy;
```

Here is an example:

Objective-C

```
id<EXBasePrx> base;
base = ...;  // Initialize base to point at a Derived
id<EXDerivedPrx> derived = [EXDerivedPrx uncheckedCast:base];
// Use derived...
```

An `uncheckedCast` provides a down-cast *without* consulting the server as to the actual run-time type of the object. You should use an `uncheckedCast` only if you are certain that the proxy indeed supports the more derived type: an `uncheckedCast`, as the name implies, is not checked in any way; it does not contact the object in the server and, if the proxy does not support the specified interface, the cast does not return null. If you use the proxy resulting from an incorrect `uncheckedCast` to invoke an operation, the behavior is undefined. Most likely, you will receive an `ICEOperationNotExistException`, but, depending on the circumstances, the Ice run time may also report an exception indicating that unmarshaling has failed, or even silently return garbage results.

Despite its dangers, `uncheckedCast` is still useful because it avoids the cost of sending a message to the server. And, particularly during initialization, it is common to receive a proxy of type `id<ICEObjectPrx>`, but with a known run-time type. In such cases, an `uncheckedCast` saves the overhead of sending a remote message.

Note that an `uncheckedCast` is *not* the same as an ordinary cast. The following is incorrect and has undefined behavior:

Objective-C

```
id<EXDerivedPrx> derived = (id<EXDerivedPrx>)base; // Wrong!
```

Both `checkedCast` and `uncheckedCast` call `autorelease` on the proxy they return so, if you want to prevent the proxy from being deallocated once the enclosing autorelease pool is drained, you must call `retain` on the returned proxy.

Using Proxy Methods in Objective-C

The `ICEObjectPrx` provides a variety of methods for customizing a proxy. Since proxies are immutable, each of these "factory methods" returns a copy of the original proxy that contains the desired modification. For example, you can obtain a proxy configured with a ten-second timeout as shown below:

Objective-C

```
id<ICEObjectPrx> proxy = [communicator stringToProxy:...];
proxy = [proxy ice_timeout:10000];
```

A factory method returns a new (autoreleased) proxy object if the requested modification differs from the current proxy, otherwise it returns the original proxy. The returned proxy is always of the same type as the source proxy, except for the factory methods `ice_facet` and `ice_identity`. Calls to either of these methods may produce a proxy for an object of an unrelated type, therefore they return a base proxy that you must subsequently down-cast to an appropriate type.

Object Identity and Proxy Comparison in Objective-C

Copyright 2016, ZeroC, Inc.

Proxy objects support comparison with `isEqual`. Note that `isEqual` uses *all* of the information in a proxy for the comparison. This means that not only the object identity must match for a comparison to succeed, but other details inside the proxy, such as the protocol and endpoint information, must be the same as well. In other words, comparison with `isEqual` tests for proxy identity, not object identity. A common mistake is to write code along the following lines:

Objective-C

```
id<ICEObjectPrx> p1 = ...;      // Get a proxy...
id<ICEObjectPrx> p2 = ...;      // Get another proxy...

if (![p1 isEqual:p2]) {
    // p1 and p2 denote different objects       // WRONG!
} else {
    // p1 and p2 denote the same object         // Correct
}
```

Even though `p1` and `p2` differ, they may denote the same Ice object. This can happen if, for example, `p1` and `p2` embed the same object identity, but use a different protocol to contact the target object. Similarly, the protocols might be the same, but could denote different endpoints (because a single Ice object can be contacted via several different transport endpoints). In other words, if two proxies compare equal with `isEqual`, we know that the two proxies denote the same object (because they are identical in all respects); however, if two proxies compare unequal with `isEqual`, we know absolutely nothing: the proxies may or may not denote the same object.

To compare the object identities of two proxies, you can use additional methods provided by proxies:

Objective-C

```
@protocol ICEObjectPrx <NSObject, NSCopying>
// ...
-(NSComparisonResult) compareIdentity:(id<ICEObjectPrx>)p;
-(NSComparisonResult) compareIdentityAndFacet:(id<ICEObjectPrx>)p;
@end
```

The `compareIdentity` method compares the object identities embedded in two proxies while ignoring other information, such as facet and transport information. To include the facet name in the comparison, use `compareIdentityAndFacet` instead.

`compareIdentity` and `compareIdentityAndFacet` allow you to correctly compare proxies for object identity. The example below demonstrates how to use `compareIdentity`:

Objective-C

```
id<ICEObjectPrx> p1 = ...;      // Get a proxy...
id<ICEObjectPrx> p2 = ...;      // Get another proxy...

if ([p1 compareIdentity:p2] != NSOrderedSame) {
    // p1 and p2 denote different objects       // Correct
} else {
    // p1 and p2 denote the same object         // Correct
}
```

See Also

- Interfaces, Operations, and Exceptions
- Proxies for Ice Objects
- Objective-C Mapping for Operations

Copyright 2016, ZeroC, Inc.

- Operations on Object
- Proxy Methods
- Versioning
- IceStorm

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Operations

On this page:

- Basic Objective-C Mapping for Operations
- Normal and idempotent Operations in Objective-C
- Passing Parameters in Objective-C
  - In-Parameters in Objective-C
  - Passing nil and NSNull in Objective-C
  - Out-Parameters in Objective-C
  - Memory Management for Out-Parameters in Objective-C
  - Receiving Return Values in Objective-C
  - Chained Invocations in Objective-C
  - nil Out-Parameters and Return Values in Objective-C
  - Optional Parameters in Objective-C
- Exception Handling in Objective-C
  - Exceptions and Out-Parameters in Objective-C
  - Exceptions and Return Values in Objective-C

Basic Objective-C Mapping for Operations

As we saw in the mapping for interfaces, for each operation on an interface, the proxy protocol contains two corresponding methods with the same name as the operation.

To invoke an operation, you call it via the proxy object. For example, here is part of the definitions for our file system:

Slice

```
["objc:prefix:FS"]
module Filesystem {
    interface Node {
        idempotent string name();
    };
    // ...
};
```

The proxy protocol for the `Node` interface looks as follows:

Objective-C

```
@protocol FSNodePrx <ICEObjectPrx>
-(NSMutableString *) name;
-(NSMutableString *) name:(ICEContext *)context;
@end;
```

The `name` method returns a value of type `NSMutableString`. Given a proxy to an object of type `Node`, the client can invoke the operation as follows:

Copyright 2016, ZeroC, Inc.

Objective-C

```
id<EXNodePrx> node = ...;       // Initialize proxy
NSString *name = [node name];   // Get name via RPC
```

The `name` method sends the operation invocation to the server, waits until the operation is complete, and then unmarshals the return value and returns it to the caller.

Because the `name` method autoreleases the return value, it is safe to ignore the return value. For example, the following code contains no memory leak:

Objective-C

```
id<EXNodePrx> node = ...;       // Initialize proxy
[node name];                    // Useless, but no leak
```

If you ignore the return value, no memory leak occurs because the next time the enclosing autorelease pool is drained, the memory will be reclaimed.

Normal and `idempotent` Operations in Objective-C

You can add an `idempotent` qualifier to a Slice operation. As far as the corresponding proxy protocol methods are concerned, `idempotent` has no effect. For example, consider the following interface:

Slice

```
interface Ops {
                string op1();
    idempotent  string op2();
    idempotent  void op3(string s);
};
```

The proxy protocol for this interface looks like this:

Objective-C

```
@protocol EXOpsPrx <ICEObjectPrx>
-(NSMutableString *) op1;
-(NSMutableString *) op1:(ICEContext *)context;
-(NSMutableString *) op2;
-(NSMutableString *) op2:(ICEContext *)context;
-(void) op3:(NSString *)s;
-(void) op3:(NSString *)s context:(ICEContext *)context;
@end;
```


Copyright 2016, ZeroC, Inc.

For brevity, we will not show the methods with the additional trailing `context` parameter for the remainder of this discussion. Of course, the compiler generates the additional methods regardless.

Because `idempotent` affects an aspect of call dispatch, not interface, it makes sense for the mapping to be unaffected by the `idempotent` keyword.

Passing Parameters in Objective-C

In-Parameters in Objective-C

The parameter passing rules for the Objective-C mapping are very simple: value type parameters are passed by value and non-value type parameters are passed by pointer. Semantically, the two ways of passing parameters are identical: the Ice run time guarantees not to change the value of an in-parameter.

Here is an interface with operations that pass parameters of various types from client to server:

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ClientToServer {
    void op1(int i, float f, bool b, string s);
    void op2(NumberAndString ns, StringSeq ss, StringTable st);
    void op3(ClientToServer* proxy);
};
```

The Slice compiler generates the following code for this definition:

Copyright 2016, ZeroC, Inc.

Objective-C

```
@interface EXNumberAndString : NSObject <NSCopying>
// ...
@property(nonatomic, assign) ICEInt x;
@property(nonatomic, retain) NSString *str;
// ...
@end

typedef NSArray EXStringSeq;
typedef NSMutableArray EXMutableStringSeq;

typedef NSDictionary EXStringTable;
typedef NSMutableDictionary EXMutableStringTable;

@protocol EXClientToServerPrx <ICEObjectPrx>
-(void) op1:(ICEInt)i f:(ICEFloat)f b:(BOOL)b s:(NSString *)s;
-(void) op2:(EXNumberAndString *)ns ss:(EXStringSeq *)ss st:(NSDictionary *)st;
-(void) op3:(id<EXClientToServerPrx>)proxy;
@end;
```

Given a proxy to a `ClientToServer` interface, the client code can pass parameters as in the following example:

Objective-C

```
id<EXClientToServerPrx> p = ...;           // Get proxy...

[p op1:42 f:3.14 b:YES s:@"Hello world!"]; // Pass literals

ICEInt i = 42;
ICEFloat f = 3.14;
BOOL b = YES;
NSString *s = @"Hello world!";

[p op1:i f:f b:b s:s];                     // Pass simple vars

EXNumberAndString *ns = [EXNumberAndString numberAndString:42 str:@"The Answer"];
EXMutableStringSeq *ss = [ExMutableStringSeq array];
[ss addObject:@"Hello world!"];
EXStringTable *st = [EXMutableStringTable dictionary];
[ss setObject:ss forKey:[NSNumber numberWithInt:0]];

[p op2:ns ss:ss st:st];                    // Pass complex vars

[p op3:p];                                 // Pass proxy
```

You can pass either literals or variables to the various operations. The Ice run time simply marshals the value of the parameters to the server and leaves parameters otherwise untouched, so there are no memory-management issues to consider.

Note that the invocation of `op3` is somewhat unusual: the caller passes the proxy it uses to invoke the operation to the operation as a

Copyright 2016, ZeroC, Inc.

parameter. While unusual, this is legal (and no memory management issues arise from doing this.)

Passing `nil` and `NSNull` in Objective-C

The Slice language supports the concept of null ("points nowhere") for only two of its types: proxies and classes. For either type, `nil` repres ents a null proxy or class. For other Slice types, such as strings, the concept of a null string simply does not apply. (There is no such thing as a null string, only the empty string.) However, strings, structures, sequences, and dictionaries are all passed by pointer, which raises the question of how the Objective-C mapping deals with `nil` values.

As a convenience feature, the Objective-C mapping permits passing of `nil` as a parameter for the following types:

- Proxies (`nil` sends a null proxy.)
- Classes (`nil` sends a null class instance.)
- Strings (`nil` sends an empty string.)
- Structures (`nil` sends a default-initialized structure.)
- Sequences (`nil` sends an empty sequence.)
- Dictionaries (`nil` sends an empty dictionary.)

It is impossible to add `nil` to an `NSArray` or `NSDictionary`, so the mapping follows the usual convention that an `NSArray` element or `NS Dictionary` value that is conceptually `nil` is represented by `NSNull`. For example, to send a sequence of proxies, some of which are null proxies, you must insert `NSNull` values into the sequence.

As a convenience feature, if you have a sequence with elements of type string, structure, sequence, or dictionary, you can use `NSNull` as the element value. For elements that are `NSNull`, the Ice run time marshals an empty string, default-initialized structure, empty sequence, or empty dictionary to the receiver.

Similarly, for dictionaries with value type string, structure, sequence, or dictionary, you can use `NSNull` as the value to send the corresponding empty value (or default-initialized value, in the case of structures).

Out-Parameters in Objective-C

The Objective-C mapping passes out-parameters by pointer (for value types) and by pointer-to-pointer (for non-value types). Here is the Slic e definition once more, modified to pass all parameters in the `out` direction:

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ServerToClient {
    void op1(out int i, out float f, out bool b, out string s);
    void op2(out NumberAndString ns,
             out StringSeq ss,
             out StringTable st);
    void op3(out ClientToServer* proxy);
};
```

The Slice compiler generates the following code for this definition:

Copyright 2016, ZeroC, Inc.

Objective-C

```
@protocol EXServerToClientPrx <ICEObjectPrx>
-(void) op1:(ICEInt *)i f:(ICEFloat *)f b:(BOOL *)b s:(NSMutableString **)s;
-(void) op2:(EXNumberAndString **)ns ss:(EXMutableStringSeq **)ss
           st:(EXMutableStringTable **)st;
-(void) op3:(id<EXClientToServerPrx> *)proxy;
@end
```

Note that, for types that come in immutable and mutable variants (strings, sequences, and dictionaries), the corresponding out-parameter uses the mutable variant.

Given a proxy to a `ServerToClient` interface, the client code can pass parameters as in the following example:

Objective-C

```
id<EXServerToClientPrx> p = ...; // Get proxy...

ICEInt i;
ICEFloat f;
BOOL b;
NSMutableString *s;

[p op1:&i f:&f b:&b s:&s];
// i, f, b, and s contain updated values now

EXNumberAndString *ns;
EXStringSeq *ss;
EXStringTable *st;

[p op2:&ns ss:&ss st:&st];
// ns, ss, and st contain updated values now

[p op3:&p];
// p has changed now!
```

Again, there are no surprises in this code: the caller simply passes pointers to pointer variables to a method; once the operation completes, the values of those variables will have been set by the server.

Memory Management for Out-Parameters in Objective-C

When the Ice run time returns an out-parameter to the caller, it does not make any assumptions about the previous value of that parameter (if any). In other words, if you pass an initialized string as an out-parameter, the value you pass is simply discarded and the corresponding variable is assigned a new instance. As an example, consider the following operation:

Slice

```
void getString(out string s);
```


Copyright 2016, ZeroC, Inc.

You could call this as follows:

Objective-C

```
NSMutableString *s = @"Hello";
[p getString:&s];
// s now points at the returned string.
```

All out-parameters are autoreleased by the Ice run time before they are returned. This is convenient because it does just the right thing with respect to memory management. For example, the following code does not leak memory:

Objective-C

```
NSMutableString *s = @"Hello";
[p getString:&s];
[p getString:&s]; // No leak here.
```

However, because the pointer value of out-parameters is simply assigned by the proxy method, you must be careful not to pass a variable as an out-parameter if that variable was not released or autoreleased:

Objective-C

```
NSMutableString *s = [[NSMutableString alloc] initWithString:@"Hello"];
[p getString:&s]; // Bad news!
```

This code leaks the initial string because the proxy method assigns the passed pointer without calling `release` on it first. (In practice, this is rarely a problem because there is no need to initialize out-parameters and, if an out-parameter was initialized by being passed as an out-parameter to an operation earlier, its value will have been autoreleased by the proxy method already.)

It is worth having another look at the final call of the code example we saw earlier:

Objective-C

```
[p op3:&p];
```

Here, p is the proxy that is used to dispatch the call. That same variable p is also passed as an out-parameter to the call, meaning that the server will set its value. In general, passing the same parameter as both an input and output parameter is safe (with the caveat we just discussed).

Receiving Return Values in Objective-C

The Objective-C mapping returns return values in much the same way as out-parameters: value types are returned by value, and non-value types are returned by pointer. As an example, consider the following operations:

Copyright 2016, ZeroC, Inc.

Slice

```
struct NumberAndString {
    int x;
    string str;
};

interface Ops {
    int getInt();
    string getString();
    NumberAndString getNumberAndString();
};
```

The proxy protocol looks as follows:

Objective-C

```
@protocol EXOpsPrx <ICEObjectPrx>
-(ICEInt) getInt;
-(NSMutableString *) getString;
-(EXNumberAndString *) getNumberAndString;
@end
```

Note that, for types with mutable and immutable variants (strings, sequences, and dictionaries), the formal return type is the mutable variant. As for out-parameters, anything returned by pointer is autoreleased by the Ice run time. This means that the following code works fine and does not contain memory management errors:

Objective-C

```
EXNumberAndString *ns = [NSNumberAndString numberAndString];
ns.x = [p getInt];
ns.str = [p getString]; // Autoreleased by getString,
                        // retained by ns.str.

[p getNumberAndString]; // No leak here.
```

The return value of `getString` is autoreleased by the proxy method but, during the assignment to the property `str`, the generated code calls `retain`, so the structure keeps the returned string alive in memory, as it should. Similarly, ignoring the return value from an invocation is safe because the returned value is autoreleased and will be reclaimed when the enclosing autorelease pool is drained.

Chained Invocations in Objective-C

Consider the following simple interface containing two operations, one to set a value and one to get it:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Name {
    string getName();
    void setName(string name);
};
```

Suppose we have two proxies to interfaces of type `Name`, `p1` and `p2`, and chain invocations as follows:

Objective-C

```
[p2 setName:[p1 getName]]; // No leak here.
```

This works exactly as intended: the value returned by `p1` is transferred to `p2`. There are no memory-management or exception safety issues with this code.

`nil` Out-Parameters and Return Values in Objective-C

If an out-parameter or return value is a proxy or class, and the operation returns a null proxy or class, the proxy method returns `nil`. If a proxy or class is returned as part of a sequence or dictionary, the corresponding entry is `NSNull`.

For strings, structures, sequences, and dictionaries, the Ice run time *never* returns `nil` or `NSNull` (even if the server passed `nil` or `NSNull` as the value). Instead, the unmarshaling code always instantiates an empty string, empty sequence, or empty dictionary, and it always initializes structure members during unmarshaling, so structures that are returned from an operation invocation never contain a `nil` instance variable (except for proxy and class instance variables).

Optional Parameters in Objective-C

The Objective-C mapping uses the `id` type for optional parameters. As a result, there's no compile time check for optional parameters. Instead, Ice performs a run-time type check and if the optional parameter does not match the expected type an `NSException` with the `NSInvalidArgumentException` name is raised. Slice types that map to an Objective-C class use the same mapping as required parameters. Slice built-in basic types (except string) are boxed into an `NSNumber` value. The `ICENone` singleton value can also be passed as the value of an optional parameter or return value.

Consider the following operation:

Slice

```
optional(1) int execute(optional(2) string aString, optional(3) int anInt, out
optional(4) float outFloat);
```

A client can invoke this operation as shown below:

Copyright 2016, ZeroC, Inc.

Objective-C

```
id f;
id i;

i = [proxy execute:@"--file log.txt" anInt:@14 outFloat:&f]
i = [proxy execute:ICENone anInt:@14 outFloat:&f] // aString is unset

if(i == ICENone)
{
    NSLog(@"return value is not set");
}
else
{
    int v = [i intValue];
    NSLog(@"return value is set to %d", v);
}
```

Passing `nil` for an optional parameter is the same as passing `nil` for a required parameter, the optional parameter is considered to be set to `nil`. For Slice built-in basic types (except string), the optional parameter is considered to be set to `0` or `NO` for booleans.

A well-behaved program must not assume that an optional parameter always has a value. It should compare the value to `ICENone` to determine whether the optional parameter is set.

Exception Handling in Objective-C

Any operation invocation may throw a run-time exception and, if the operation has an exception specification, may also throw user exceptions. Suppose we have the following simple interface:

Slice

```
exception Tantrum {
    string reason;
};

interface Child {
    void askToCleanUp() throws Tantrum;
};
```

Slice exceptions are thrown as Objective-C exceptions, so you can simply enclose one or more operation invocations in a `try-catch` block:

Copyright 2016, ZeroC, Inc.

Objective-C

```
id<EXChildPrx> child = ...;    // Get proxy...
@try {
    [child askToCleanUp];      // Give it a try...
} @catch (EXTantrum *t) {
    printf("The child says: %s\n", t.reason_);
}
```

Typically, you will catch only a few exceptions of specific interest around an operation invocation; other exceptions, such as unexpected run-time errors, will typically be dealt with by exception handlers higher in the hierarchy. For example:

Objective-C

```
void run()
{
    id<EXChildPrx> child = ...;       // Get proxy...
    @try {
        [child askToCleanUp];         // Give it a try...
    } @catch (EXTantrum *t) {
        printf("The child says: %s\n", t.reason);
        [child scold];                // Recover from error...
    }
    [child praise];                   // Give positive feedback...
}

int
main(int argc, char* argv[])
{
    int status = 1;
    @try {
        // ...
        run();
        // ...
        status = 0;
    } @catch (ICEException *e) {
        printf("Unexpected run-time error: %s\n", [e ice_name]);
    }
    // ...
    return status;
}
```

This code handles a specific exception of local interest at the point of call and deals with other exceptions generically. (This is also the strategy we used for our first simple application.)

Exceptions and Out-Parameters in Objective-C

If an operation throws an exception, the Ice run time makes no guarantee for the value of out-parameters. Individual out-parameters may have the old value, the new value, or a value that is indeterminate, such that parts of the out-parameter have been assigned and others have not. However, no matter what their state, the values will be "safe" for memory-management purposes, that is, any out-parameters that were successfully unmarshaled are autoreleased.

Copyright 2016, ZeroC, Inc.

Exceptions and Return Values in Objective-C

For return values, the Objective-C mapping provides the guarantee that a variable receiving the return value of an operation will not be overwritten if an exception is thrown.

See Also

- Operations
- Hello World Application
- Objective-C Mapping for Interfaces
- Objective-C Mapping for Exceptions
- Objective-C Mapping for Local Interfaces
- Request Contexts

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Local Interfaces

The Ice run time defines APIs using Slice. These APIs are provided as part of the Ice run time library and cannot be invoked remotely. (Doing so would not make any sense.) Therefore, the Slice interfaces for the Ice run time are defined as local interfaces. The Objective-C mapping for local interfaces differs from the default mapping in two ways:

- Local interfaces do not adopt an `<`*`interface-name`*`>Prx` protocol. (Doing so would be misleading because proxies imply that the target object can be remote.) Instead, the protocol for local interfaces has the same name as the interface. For example, the `Ice::Communicator` interfaces is defined as:

  Slice

  ```
  ["objc:prefix:ICE"]
  module Ice {
      local interface Communicator {
          // ...
      }
  };
  ```

  Because `Communicator` is a local interface, objects of type `ICECommunicator` are passed as `id<ICECommunicator>` (*not* `ICECommunicator*` or `id<ICEComunicatorPrx>`).

- Types that come in mutable and immutable variants (strings, sequences, and dictionaries) are always passed as the immutable variant. For example, the `getName` operation on the `ObjectAdapter` interface is defined as:

  Slice

  ```
  ["objc:prefix:ICE"]
  module Ice {
      local interface ObjectAdapter {
          string getName();
      };
  };
  ```

  Because `ObjectAdapter` is a local interface, the `getName` operation maps to:

  Objective-C

  ```
  -(NSString *) getName;
  ```

  Note that the returned string is of type `NSString` instead of `NSMutableString` (as would be the case for an operation on a non-local interface).
  For local interfaces, parameters are passed as the immutable version because their values are not meant to be modified by application code. In addition, passing the immutable version avoids an unnecessary data copy.

See Also

- Modules
- Local Types
- Objective-C Mapping for Operations

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Classes

On this page:

- Basic Objective-C Mapping for Classes
- Derivation from ICEObject in Objective-C
- Class Data Members in Objective-C
- Class Constructors in Objective-C
- Derived Classes in Objective-C
- Passing Classes as Parameters in Objective-C
- Operations of Classes in Objective-C
- Class Factories in Objective-C
    - Using a Category to Implement Operations in Objective-C
- Copying of Classes in Objective-C
    - Cyclic References in Objective-C

Basic Objective-C Mapping for Classes

A Slice class is mapped similar to a structure and exception.

The generated class contains an instance variable and a property for each Slice data member. Consider the following class definition:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Slice compiler generates the following code for this definition:

Objective-C

```
@interface EXTimeOfDay : ICEObject
{
    ICEShort hour;
    ICEShort minute;
    ICEShort second;
}

@property(nonatomic, assign) ICEShort hour;
@property(nonatomic, assign) ICEShort minute;
@property(nonatomic, assign) ICEShort second;

-(id) init:(ICEShort)hour minute:(ICEShort)minute second:(ICEShort)second;
+(id) timeOfDay;
+(id) timeOfDay:(ICEShort)hour minute:(ICEShort)minute second:(ICEShort)second;
@end
```

There are a number of things to note about the generated code:

1. The generated class `EXTimeOfDay` derives from `ICEObject`, which is the parent of all classes. Note that `ICEObject` is *not* the

Copyright 2016, ZeroC, Inc.

same as `ICEObjectPrx`. In other words, you *cannot* pass a class where a proxy is expected and vice versa.
2. The generated class contains a property for each Slice data member.
3. The generated class provides an `init` method that accepts one argument for each data member, and it provides the same two convenience constructors as structures and exceptions.

Derivation from `ICEObject` in Objective-C

All classes ultimately derive from a common base class, `ICEObject`. Note that this is not the same as implementing the `ICEObjectPrx` protocol (which is implemented by proxies). As a result, you cannot pass a class where a proxy is expected (and vice versa) because the base types for classes and proxies are not compatible.

`ICEObject` defines a number of methods:

Objective-C

```
@protocol ICEObject <NSObject>
-(BOOL) ice_isA:(NSString*)typeId current:(ICECurrent*)current;
-(void) ice_ping:(ICECurrent*)current;
-(NSString*) ice_id:(ICECurrent*)current;
-(NSArray*) ice_ids:(ICECurrent*)current;
@end

@interface ICEObject NSObject <ICEObject, NSCopying>
-(BOOL) ice_isA:(NSString*)typeId;
-(void) ice_ping;
-(NSString*) ice_id;
-(NSArray*) ice_ids;
+(NSString*) ice_staticId;
-(void) ice_preMarshal;
-(void) ice_postUnmarshal;
-(BOOL) ice_dispatch:(id<ICERequest>)request;
-(id) initWithDelegate:(id)delegate;
+(id) objectWithDelegate:(id)delegate;
@end
```


The methods are split between the `ICEObject` protocol and class because classes can be servants.

The methods of `ICEObject` behave as follows:

- `ice_isA`
  This function returns `YES` if the object supports the given type ID, and `NO` otherwise.
- `ice_ping`
  `ice_ping` provides a basic reachability test for the class. If it completes without raising an exception, the class exists and is reachable. Note that `ice_ping` is normally only invoked on the proxy for a class that might be remote because a class instance that is local (in the caller's address space) can always be reached.
- `ice_ids`
  This function returns a string sequence representing all of the type IDs supported by this object, including `::Ice::Object`.
- `ice_id`
  This function returns the actual run-time type ID for a class. If you call `ice_id` via a pointer to a base instance, the returned type ID is the actual (possibly more derived) type ID of the instance.
- `ice_staticId`
  This function returns the static type ID of a class.
- `ice_preMarshal`

Copyright 2016, ZeroC, Inc.

The Ice run time invokes this function prior to marshaling the object's state, providing the opportunity for a subclass to validate its declared data members.

- `ice_postUnmarshal`
  The Ice run time invokes this function after unmarshaling an object's state. A subclass typically overrides this function when it needs to perform additional initialization using the values of its declared data members.
- `ice_dispatch`
  This function dispatches an incoming request to a servant. It is used in the implementation of dispatch interceptors.
- `initWithDelegate`
  These constructors enable the implementation of servants with a delegate.

Class Data Members in Objective-C

By default, data members of classes are mapped exactly as for structures and exceptions: for each data member in the Slice definition, the generated class contains a corresponding property.

Class Constructors in Objective-C

Classes provide the usual `init` method and a parameter-less convenience constructor that perform default initialization of the object's instance variables. These methods initialize the instance variables with zero-filled memory, with the following exceptions:

- A string data member is initialized to an empty string
- An enumerator data member is initialized to the first enumerator
- A structure data member is initialized to a default-constructed value

If you declare default values in your Slice definition, the `init` method and convenience constructor initialize each data member with its declared value instead.

In addition, if a class has data members, it provides an `init` method and a convenience constructor that accept one argument for each data member. This allows you to allocate and initialize a class instance in a single statement (instead of first having to allocate and default-initialize the instance and then assign to its properties).

For derived classes, the `init` method and the convenience constructor have one parameter for each of the base class's data members, plus one parameter for each of the derived class's data members, in base-to-derived order. For example:

Slice

```
class Base {
    int i;
};

class Derived extends Base {
    string s;
};
```

This generates:

Copyright 2016, ZeroC, Inc.

Objective-C

```
@interface EXBase : ICEObject
// ...

@property(nonatomic, assign) ICEInt i;

-(id) init:(ICEInt)i;
+(id) base;
+(id) base:(ICEInt)i;
@end

@interface EXDerived : EXBase
// ...

@property(nonatomic, retain) NSString *s;

-(id) init:(ICEInt)i s:(NSString *)s;
+(id) derived;
+(id) derived:(ICEInt)i s:(NSString *)s;
@end
```

Derived Classes in Objective-C

Note that, in the preceding example, the derivation of the Slice definitions is preserved for the generated classes: `EXBase` derives from `ICEObject`, and `EXDerived` derives from `EXBase`. This allows you to treat and pass classes polymorphically: you can always pass an `EXDerived` instance where an `EXBase` instance is expected.

Passing Classes as Parameters in Objective-C

Classes are passed by pointer, like any other Objective-C object. For example, here is an operation that accepts a `Base` as an in-parameter and returns a `Derived`:

Slice

```
Derived getDerived(Base d);
```

The corresponding proxy method looks as follows:

Objective-C

```
-(EXDerived *) getDerived:(EXBase *)d;
```

To pass a null instance, you simply pass `nil`.

Operations of Classes in Objective-C

Copyright 2016, ZeroC, Inc.

If you look back at the code that is generated for the `EXTimeOfDay` class, you will notice that there is no indication at all that the class has a `format` operation. As opposed to proxies, classes do not implement any protocol that would define which operations are available. This means that you can partially implement the operations of a class. For example, you might have a Slice class with five operations that is returned from a server to a client. If the client uses only one of the five operations, the client-side code needs to implement only that one operation and can leave the remaining four operations without implementation. (If the class were to implement a mandatory protocol, the client-side code would have to implement all operations in order to avoid a compiler warning.)

Of course, you must implement those operations that you actually intend to call. The mapping of operations for classes follows the *server-side* mapping for operations on interfaces: parameter types and labels are exactly the same. (See Parameter Passing in Objective-C for details.) In a nutshell, the server-side mapping is the same as the client-side mapping except that, for types that have mutable and immutable variants, they map to the immutable variant where the client-side mapping uses the mutable variant, and vice versa.

For example, here is how we could implement the `format` operation of our `TimeOfDay` class:

Objective-C

```
@interface TimeOfDayI : EXTimeOfDay
@end

@implementation TimeOfDayI
-(NSString *) format
{
    return [NSString stringWithFormat:@"%.2d:%.2d:%.2d",
self.hour, self.minute, self.second];
}
@end
```

By convention, the implementation of classes with operations has the same name as the Slice class with an `I`-suffix. Doing this is not mandatory — you can call your implementation class anything you like. However, if you do not want to use the `I`-suffix naming, we recommend that you adopt another naming convention and follow it consistently.

Note that `TimeOfDayI` derives from `EXTimeOfDay`. This is because, as we will see in a moment, the Ice run time will instantiate a `TimeOfDayI` instance whenever it receives a `TimeOfDay` instance over the wire and expects that instance to provide the properties of `EXTimeOfDay`.

Class Factories in Objective-C

Having created a class such as `TimeOfDayI`, we have an implementation and we can instantiate the `TimeOfDayI` class, but we cannot receive it as the return value or as an out-parameter from an operation invocation. To see why, consider the following simple interface:

Slice

```
interface Time {
    TimeOfDay get();
};
```

When a client invokes the `get` operation, the Ice run time must instantiate and return an instance of the `TimeOfDayI` class. However, unless we tell it, the Ice run time cannot magically know that we have created a `TimeOfDayI` class that implements a `format` method. To allow the Ice run time to instantiate the correct object, we must provide a factory that knows that the Slice `TimeOfDay` class is implemented by our `TimeOfDayI` class. The `Ice::Communicator` interface provides us with the necessary operations:

Copyright 2016, ZeroC, Inc.

Slice

```
["objc:prefix:ICE"]
module Ice {
    local interface ObjectFactory {
        Object create(string type);
        void destroy();
    };

    local interface Communicator {
        void addObjectFactory(ObjectFactory factory, string id);
        ObjectFactory findObjectFactory(string id);
        // ...
    };
};
```

To supply the Ice run time with a factory for our `TimeOfDayI` class, we must implement the `ObjectFactory` interface:

Slice

```
module Ice {
    local interface ObjectFactory {
        Object create(string type);
        void destroy();
    };
};
```

The object factory's `create` operation is called by the Ice run time when it needs to instantiate a `TimeOfDay` class. The factory's `destroy` operation is called by the Ice run time when its communicator is destroyed. A possible implementation of our object factory is:

Objective-C

```
@interface ObjectFactory<ICEObjectFactory>
@end

@implementation ObjectFactory
-(ICEObject*) create:(NSString *)type
{
    NSAssert([type isEqualToString:@"::Example::TimeOfDay"]);
    return [[TimeOfDayI alloc] init];
}
@end
```

The `create` method is passed the type ID of the class to instantiate. For our `TimeOfDay` class, the type ID is `"::Example::TimeOfDay"`. Our implementation of `create` checks the type ID: if it is `"::Example::TimeOfDay"`, it instantiates and returns a `TimeOfDayI` object. For other type IDs, it asserts because it does not know how to instantiate other types of objects.

Copyright 2016, ZeroC, Inc.

Note that your factory *must not* autorelease the returned instance. The Ice run time takes care of the necessary memory management activities on your behalf.

Given a factory implementation, such as our `ObjectFactory`, we must inform the Ice run time of the existence of the factory:

Objective-C

```
id<ICECommunicator> ice = ...;
ObjectFactory *factory = [[[ObjectFactory alloc] init] autorelease];
[ic addObjectFactory:factory sliceId:@"::Example::TimeOfDay"];
```

Now, whenever the Ice run time needs to instantiate a class with the type ID `"::Example::TimeOfDay"`, it calls the `create` method of the registered `ObjectFactory` instance.

The `destroy` operation of the object factory is invoked by the Ice run time when the communicator is destroyed. This gives you a chance to clean up any resources that may be used by your factory. Do not call `destroy` on the factory while it is registered with the communicator — if you do, the Ice run time has no idea that this has happened and, depending on what your `destroy` implementation is doing, may cause undefined behavior when the Ice run time tries to next use the factory.

The run time guarantees that `destroy` will be the last call made on the factory, that is, `create` will not be called concurrently with `destroy`, and `create` will not be called once `destroy` has been called. However, the Ice run time may make concurrent calls to `create`.

Note that you cannot register a factory for the same type ID twice: if you call `addObjectFactory` with a type ID for which a factory is registered, the Ice run time throws an `AlreadyRegisteredException`.

Finally, keep in mind that if a class has only data members, but no operations, you need not (but can) create and register an object factory to transmit instances of such a class. Only if a class has operations do you have to define and register an object factory.

Using a Category to Implement Operations in Objective-C

An alternative to registering a class factory is to use an Objective-C category to implement operations. For example, we could have implemented our `format` method using a category instead:

Objective-C

```
@interface EXTimeOfDay (TimeOfDayI)
@end

@implementation EXTimeOfDay (TimeOfDayI)
-(NSString *) format
{
    return [NSString stringWithFormat:@"%.2d:%.2d:%.2d",
self.hour, self.minute, self.second];
}
@end
```

In this case, there is no need to derive from the generated `EXTimeOfDay` class because we provide the format implementation as a category. There is also no need to register a class factory: the Ice run time instantiates an `EXTimeOfDay` instance when a `TimeOfDay` instance arrives over the wire, and the `format` method is found at run time when it is actually called.

This is a viable alternative approach to implement class operations. However, keep in mind that, if the operation implementation requires use of instance variables that are not defined as part of the Slice definitions of a class, you cannot use this approach because Objective-C categories do not permit you to add instance variables to a class.

Copying of Classes in Objective-C

Copyright 2016, ZeroC, Inc.

Classes implement `NSCopying`. The behavior is the same as for structures: instance variables of value type are copied by assignment, instance variables of pointer type are copied by calling `retain`, that is, the copy is shallow. To illustrate this, consider the following class definition:

Slice

```
class Node {
    int i;
    string s;
    Node next;
};
```

We can initialize two instances of type `EXNode` as follows:

Objective-C

```
NSString lastString = [NSString stringWithString:@"last"];
EXNode *last = [EXNode node:99 s:lastString next:nil];

NSString firstString = [NSString stringWithString:@"first"];
EXNode *first = [EXNode node:1 s:firstString next:last];
```

This creates the situation shown below:

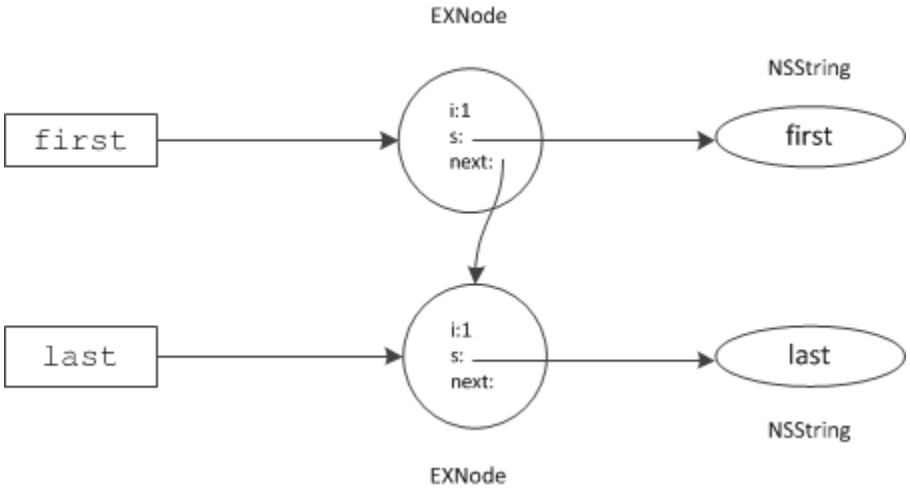

Two instances of type EXNode.

Now we create a copy of the first node by calling `copy`:

Objective-C

```
EXNode *copy = [[first copy] autorelease];
```

This creates the situation shown here:

Copyright 2016, ZeroC, Inc.

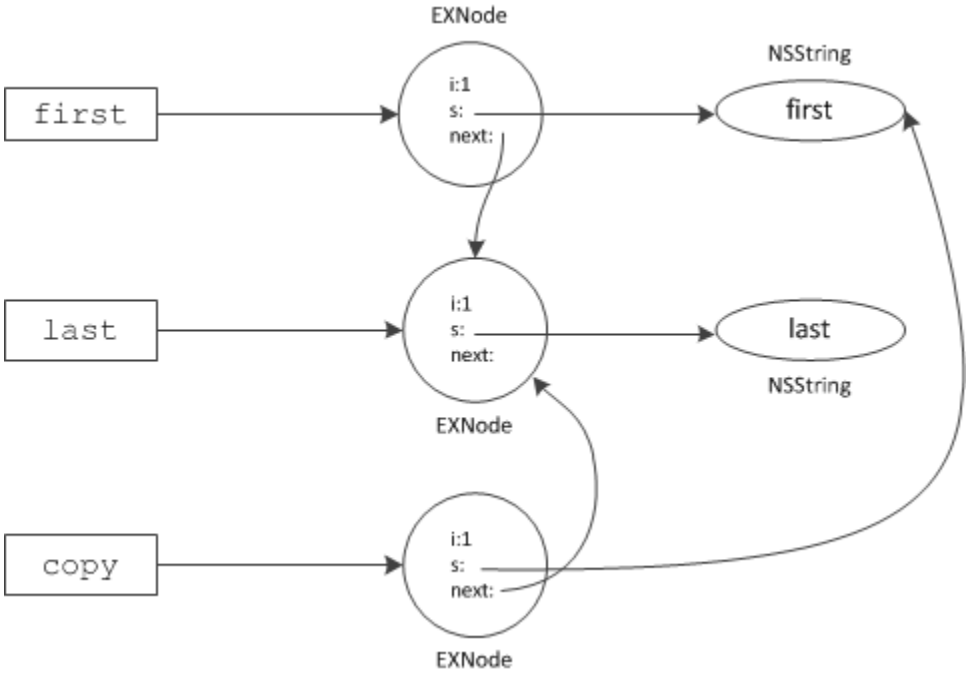

EXNode instances after calling copy on first.

As you can see, the first node is copied, but the last node (pointed at by the `next` instance variable of the first node) is not copied; instead, `first` and `copy` now both have their `next` instance variable point at the same last node, and both point at the same string.

Cyclic References in Objective-C

One thing to be aware of are cyclic references among classes. As an example, we can easily create a cycle by executing the following statements:

Objective-C

```
EXNode *first = [EXNode node];
ExNode *last = [EXNode node];
first.next = last;
last.next = first;
```

This makes the `next` instance variable of the two classes point at each other, creating the cycle shown below:

Copyright 2016, ZeroC, Inc.

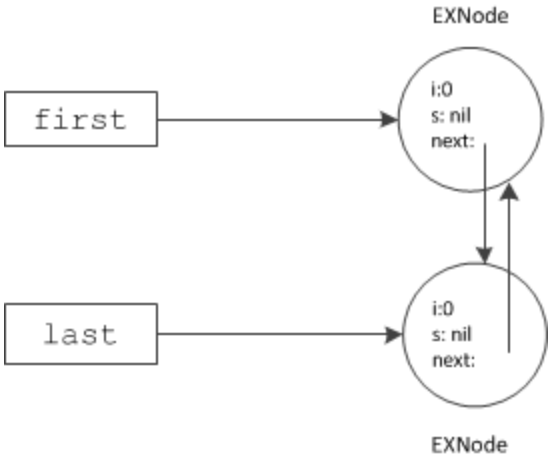

Two nodes with cyclic references.

There is no problem with sending this class graph as a parameter. For example, you could pass either `first` or `last` as a parameter to an operation and, in the server, the Ice run time will faithfully rebuild the corresponding graph, preserving the cycle. However, if a server returns such a graph from an operation invocation as the return value or as an out-parameter, all class instances that are part of a cycle are leaked. The same is true on the client side: if you receive such a graph from an operation invocation and do not explicitly break the cycle, you will leak all instances that form part of the cycle.

Because it is difficult to break cycles manually (and, on the server side, for return values and out-parameters, it is impossible to break them), we recommend that you avoid cyclic references among classes.

A future version of the Objective-C run time may provide a garbage collection facility similar to the one provided with the C++ mapping.

See Also

- Simple Classes
- Objective-C Mapping for Classes]
- Server-Side Objective-C Mapping for Interfaces
- Parameter Passing in Objective-C
- Dispatch Interceptors

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Interfaces by Value

Slice permits you to pass an interface by value:

Slice

```
interface ClassBase {
    void someOp();
    // ...
};

interface Processor {
    ClassBase process(ClassBase b);
};

class SomeClass implements ClassBase {
    // ...
};
```

Note that `process` accepts and returns a value of type `ClassBase`. This is *not* the same as passing `id<ClassBasePrx>`, which is a *proxy* to an object of type `ClassBase` that is possibly remote. Instead, what is passed here is an *interface*, and the interface is passed by *value*.

The immediate question is "what does this mean?" After all, interfaces are abstract and, therefore, it is impossible to pass an interface by value. The answer is that, while an interface cannot be passed, what *can* be passed is a class that implements the interface. That class is type-compatible with the formal parameter type and, therefore, can be passed by value. In the preceding example, `SomeClass` implements `ClassBase` and, hence, can be passed to and returned from the `process` operation.

The Objective-C mapping maps interface-by-value parameters to `ICEObject*`, regardless of the type of the interface. For example, the proxy protocol for the `process` operation is:

Objective-C

```
-(ICEObject *) process:(ICEObject *)b;
```

This means that you can pass a class of any type to the operation, even if it is not type-compatible with the formal parameter type, because all classes derive from `ICEObject`. However, an invocation of `process` is still type-safe at run time: the Ice run time verifies that the class instance that is passed implements the specified interface; if not, the invocation throws an `ICEMarshalException`.

Passing interfaces by value as `ICEObject*` is a consequence of the decision to not generate a formal protocol for classes. (If such a protocol would exist, the formal parameter type could be `id<`*`ProtocolName`*`>`. However, as we described for the class mapping, a protocol would require the implementation of a class to implement all of its operations, which can be inconvenient. Because it is rare to pass interfaces by value (more often, the formal parameter type will be a base *class* instead of a base *interface*), the minor loss of static type safety is an acceptable trade-off.

See Also

- Passing Interfaces by Value
- Objective-C Mapping for Classes

Copyright 2016, ZeroC, Inc.

Objective-C Mapping for Optional Data Members

The Objective-C mapping for optional data members in Slice classes and exceptions adds an extra boolean instance variable as well as two selectors for testing if the optional is set and clearing its value. The argument for the optional data member in the convenience constructor is mapped according to the optional parameter mapping for Slice operations. Consider the following Slice definition:

Slice

```
class C
{
    string name;
    optional(2) string alternateName;
    optional(5) bool active;
};
```

The generated Objective-C code provides the following API:

Objective-C

```
@interface testC : ICEObject
{
    NSString *name;
    NSString *alternateName;
    BOOL has_alternateName__;
    BOOL active;
    BOOL has_active__;
}
@property(nonatomic, ICE_STRONG_ATTR) NSString *name;
@property(nonatomic, ICE_STRONG_ATTR) NSString *alternateName;
@property(nonatomic, assign) BOOL active;
-(id) init:(NSString*)name alternateName:(id)alternateName active:(id)active;
+(id) c:(NSString*)name alternateName:(id)alternateName active:(id)active;
+(id) c;
-(void)setAlternateName:(NSString*)alternateName_p;
-(BOOL)hasAlternateName;
-(void)clearAlternateName;
-(void)setActive:(BOOL)active_p;
-(BOOL)hasActive;
-(void)clearActive;
@end
```

Note that calling a `get` method or accessing the property when the value is not currently set returns an undefined value.

See Also

- Optional Data Members

Copyright 2016, ZeroC, Inc.

Asynchronous Method Invocation (AMI) in Objective-C

Asynchronous Method Invocation (AMI) is the term used to describe the client-side support for the asynchronous programming model. AMI supports both oneway and twoway requests, but unlike their synchronous counterparts, AMI requests never block the calling thread. When a client issues an AMI request, the Ice run time hands the message off to the local transport buffer or, if the buffer is currently full, queues the request for later delivery. The application can then continue its activities and poll or wait for completion of the invocation, or receive a callback when the invocation completes.

AMI is transparent to the server: there is no way for the server to tell whether a client sent a request synchronously or asynchronously.

On this page:

- Basic Asynchronous API in Objective-C
  - Proxy Methods for AMI in Objective-C
  - Exception Handling for AMI in Objective-C
- The ICEAsyncResult Protocol in Objective-C
- Polling for Completion in Objective-C
- Completion Callbacks in Objective-C
- Oneway Invocations in Objective-C
- Flow Control in Objective-C
- Batch Requests in Objective-C
- Concurrency in Objective-C

Basic Asynchronous API in Objective-C

Consider the following simple Slice definition:

Slice

```
module Demo {
    interface Employees {
        string getName(int number);
    };
};
```

Proxy Methods for AMI in Objective-C

Besides the synchronous proxy methods, the Objective-C mapping generates the following asynchronous proxy methods:

Copyright 2016, ZeroC, Inc.

Objective-C

```
-(id<ICEAsyncResult>) begin_getName:(ICEInt)number;

-(id<ICEAsyncResult>) begin_getName:(ICEInt)number
context:(ICEContext *)context;

-(id<ICEAsyncResult>) begin_getName:(ICEInt)number
response:(void(^)(NSMutableString*))response_
exception:(void(^)(ICEException*))exception_;

-(id<ICEAsyncResult>) begin_getName:(ICEInt)number
context:(ICEContext *)context
response:(void(^)(NSMutableString*))response_
exception:(void(^)(ICEException*))exception_;

-(id<ICEAsyncResult>) begin_getName:(ICEInt)number
response:(void(^)(NSMutableString*))response_
exception:(void(^)(ICEException*))exception_
sent:(void(^)(BOOL))sent_;

-(id<ICEAsyncResult>) begin_getName:(ICEInt)number
context:(ICEContext *)context
response:(void(^)(NSMutableString*))response_
exception:(void(^)(ICEException*))exception_
sent:(void(^)(BOOL))sent_;

-(NSMutableString *) end_getName:(id<ICEAsyncResult>)result;
```

As you can see, the single `getName` operation results in several `begin_getName` methods as well as an `end_getName` method. The `begin_` methods optionally accept a per-invocation context and callbacks.

- The `begin_getName` methods send (or queue) an invocation of `getName`. These methods do not block the calling thread.
- The `end_getName` method collects the result of the asynchronous invocation. If, at the time the calling thread calls `end_getName`, the result is not yet available, the calling thread blocks until the invocation completes. Otherwise, if the invocation completed some time before the call to `end_getName`, the method returns immediately with the result.

A client could call these methods as follows:

Objective-C

```
id<EXEmployeesPrx> e = [EXEmployeesPrx checkedCast:...];
id<ICEAsyncResult> r = [e begin_getName:99]

// Continue to do other things here...

NSString* name = [e end_getName:r];
```

Because `begin_getName` does not block, the calling thread can do other things while the operation is in progress.

Note that `begin_getName` returns a value of type `id<ICEAsyncResult>`. This value contains the state that the Ice run time requires to

Copyright 2016, ZeroC, Inc.

keep track of the asynchronous invocation. You must pass the `id<ICEAsyncResult>` that is returned by the `begin_` method to the corresponding `end_` method.

The `begin_` method has one parameter for each in-parameter of the corresponding Slice operation. The `end_` method accepts the `id<ICEAsyncResult>` object as its only argument and returns the out-parameters using the same semantics as for regular synchronous invocations. For example, consider the following operation:

Slice

```
double op(int inp1, string inp2, out bool outp1, out long outp2);
```

The `begin_op` and `end_op` methods have the following signature:

Objective-C

```
-(id<ICEAsyncResult>) begin_op:(ICEInt)inp1 inp2:(NSString *)inp2;
-(ICEDouble) end_op:(BOOL*)outp1 outp2:(ICELong*)outp2
result:(id<ICEAsyncResult>)result;
```

The call to `end_op` returns the out-parameters as follows:

Objective-C

```
BOOL outp1;
ICELong outp2;
ICEDouble doubleValue = [p end_op:&outp1 outp2:&outp2 result:result];
```

Exception Handling for AMI in Objective-C

If an invocation raises an exception, the exception is thrown by the `end_` method, even if the actual error condition for the exception was encountered during the `begin_` method ("on the way out"). The advantage of this behavior is that all exception handling is located with the code that calls the `end_` method (instead of being present twice, once where the `begin_` method is called, and again where the `end_` method is called).

There is one exception to the above rule: if you destroy the communicator and then make an asynchronous invocation, the `begin_` method throws `ICECommunicatorDestroyedException`. This is necessary because, once the run time is finalized, it can no longer throw an exception from the `end_` method.

The only other exception that is thrown by the `begin_` and `end_` methods is `NSException` with the `NSInvalidArgumentException` name. This exception indicates that you have used the API incorrectly. For example, the `begin_` method throws this exception if you call an operation that has a return value or out-parameters on a oneway proxy. Similarly, the `end_` method throws this exception if you use a different proxy to call the `end_` method than the proxy you used to call the `begin_` method, or if the `id<ICEAsyncResult>` you pass to the `end_` method was obtained by calling the `begin_` method for a different operation.

The `ICEAsyncResult` Protocol in Objective-C

The `id<ICEAsyncResult>` that is returned by the `begin_` method encapsulates the state of the asynchronous invocation:

Copyright 2016, ZeroC, Inc.

Objective-C

```
@protocol ICEAsyncResult <NSObject>
-(id<ICECommunicator>) getCommunicator;
-(id<ICEConnection>) getConnection;
-(id<ICEObjectPrx>) getProxy;

-(BOOL) isCompleted;
-(void) waitForCompleted;

-(BOOL) isSent;
-(void) waitForSent;

-(BOOL) sentSynchronously;
-(NSString*) getOperation;
@end
```

The methods have the following semantics:

- `getCommunicator`
  This method returns the communicator that sent the invocation.
- `getConnection`
  This method returns the connection that was used for the invocation. Note that, for typical asynchronous proxy invocations, this method returns a nil value because the possibility of automatic retries means the connection that is currently in use could change unexpectedly. The `getConnection` method only returns a non-nil value when the `ICEAsyncResult` is obtained by calling `begin_flushBatchRequests` on a `Connection` object.
- `getProxy`
  This method returns the proxy that was used to call the `begin_` method, or nil if the `ICEAsyncResult` was not obtained via an asynchronous proxy invocation.
- `getOperation`
  This method returns the name of the operation.
- `isCompleted`
  This method returns true if, at the time it is called, the result of an invocation is available, indicating that a call to the `end_` method will not block the caller. Otherwise, if the result is not yet available, the method returns false.
- `waitForCompleted`
  This method blocks the caller until the result of an invocation becomes available.
- `isSent`
  When you call the `begin_` method, the Ice run time attempts to write the corresponding request to the client-side transport. If the transport cannot accept the request, the Ice run time queues the request for later transmission. `isSent` returns true if, at the time it is called, the request has been written to the local transport (whether it was initially queued or not). Otherwise, if the request is still queued, `isSent` returns false.
- `waitForSent`
  This method blocks the calling thread until a request has been written to the client-side transport.
- `sentSynchronously`
  This method returns true if a request was written to the client-side transport without first being queued. If the request was initially queued, `sentSynchronously` returns false (independent of whether the request is still in the queue or has since been written to the client-side transport).

Polling for Completion in Objective-C

The `ICEAsyncResult` methods allow you to poll for call completion. Polling is useful in a variety of cases. As an example, consider the following simple interface to transfer files from client to server:

Copyright 2016, ZeroC, Inc.

Slice

```
interface FileTransfer
{
    void send(int offset, ByteSeq bytes);
};
```

The client repeatedly calls `send` to send a chunk of the file, indicating at which offset in the file the chunk belongs. A naïve way to transmit a file would be along the following lines:

Objective-C

```
NSInputStream* stream = ...
id<EXFileTransferPrx> ft = [EXFileTransferPrx checkedCast:...];
int chunkSize = ...;
int offset = 0;
while([stream hasBytesAvailable])
{
    char bytes[chunkSize];
    int l = [stream read:bytes maxLength:sizeof(bytes)];
    if(l > 0)
    {
        [ft send:offset bytes:[ByteSeq dataWithBytes:bytes length:l]];
        offset += l;
    }
}
```

This works, but not very well: because the client makes synchronous calls, it writes each chunk on the wire and then waits for the server to receive the data, process it, and return a reply before writing the next chunk. This means that both client and server spend much of their time doing nothing — the client does nothing while the server processes the data, and the server does nothing while it waits for the client to send the next chunk.

Using asynchronous calls, we can improve on this considerably:

Copyright 2016, ZeroC, Inc.

Objective-C

```
NSInputStream* stream = ...
id<EXFileTransferPrx> ft = [EXFileTransferPrx checkedCast:...];
int chunkSize = ...;
int offset = 0;
NSMutableArray* results = [NSMutableArray arrayWithCapacity:5];
int numRequests = 5;
while([stream hasBytesAvailable])
{
    char bytes[chunkSize];
    int l = [stream read:bytes maxLength:sizeof(bytes)];
    if(l > 0)
    {
        // Send up to numRequests + 1 chunks asynchronously.
        id<ICEAsyncResult> r =
            [ft begin_send:offset bytes:[ByteSeq dataWithBytes:bytes length:l]];
        offset += l;

        // Wait until this request has been passed to the
        // transport.
        [r waitForSent];
        [results addObject:r];

        // Once there are more than numRequests, wait for the
        // least recent one to complete.
        while([results count] > numRequests)
        {
            r = [results objectAtIndex:0];
            [results removeObjectAtIndex:0];
            [r waitForCompleted];
        }
    }
}

// Wait for any remaining requests to complete.
for(id<ICEAsyncResult> r in results)
{
    [r waitForCompleted];
}
```

With this code, the client sends up to `numRequests + 1` chunks before it waits for the least recent one of these requests to complete. In other words, the client sends the next request without waiting for the preceding request to complete, up to the limit set by `numRequests`. In effect, this allows the client to "keep the pipe to the server full of data": the client keeps sending data, so both client and server continuously do work.

Obviously, the correct chunk size and value of `numRequests` depend on the bandwidth of the network as well as the amount of time taken by the server to process each request. However, with a little testing, you can quickly zoom in on the point where making the requests larger or queuing more requests no longer improves performance. With this technique, you can realize the full bandwidth of the link to within a percent or two of the theoretical bandwidth limit of a native socket connection.

Completion Callbacks in Objective-C

The `begin_` method accepts three optional callback arguments that allow you to be notified asynchronously when a request completes.

Copyright 2016, ZeroC, Inc.

Here is the signature of the `begin_getName` method that we saw earlier:

Objective-C

```
-(id<ICEAsyncResult>) begin_getName:(ICEInt)number
    response:(void(^)(NSMutableString*))response_
    exception:(void(^)(ICEException*))exception_;

-(id<ICEAsyncResult>) begin_getName:(ICEInt)number
    context:(ICEContext *)context
    response:(void(^)(NSMutableString*))response_
    exception:(void(^)(ICEException*))exception_;

-(id<ICEAsyncResult>) begin_getName:(ICEInt)number
    response:(void(^)(NSMutableString*))response_
    exception:(void(^)(ICEException*))exception_
    sent:(void(^)(BOOL))sent_;

-(id<ICEAsyncResult>) begin_getName:(ICEInt)number
    context:(ICEContext *)context
    response:(void(^)(NSMutableString*))response_
    exception:(void(^)(ICEException*))exception_
    sent:(void(^)(BOOL))sent_;
```

The value you pass for the response callback (`response`), the exception callback (`exception`), or the sent callback (`sent`) argument must be an Objective-C block. The response callback is invoked when the request completes successfully, and the exception callback is invoked when the operation raises an exception. (The sent callback is primarily used for flow control.)

For example, consider the following callbacks for an invocation of the `getName` operation:

Objective-C

```
void(^getNameCB)(NSMutableString*) = ^(NSMutableString* name)
{
    NSLog(@"Name is: %@", name);
};

void(^failureCB)(ICEException*) = ^(ICEException* ex)
{
    NSLog(@"Exception is: %@", [ex description]);
};
```

The response callback parameters depend on the operation signature. If the operation has a non-`void` return type, the first parameter of the response callback is the return value. The return value (if any) is followed by a parameter for each out-parameter of the corresponding Slice operation, in the order of declaration.

The exception callback is called if the invocation fails because of an Ice run time exception, or if the operation raises a user exception.

To inform the Ice run time that you want to receive callbacks for the completion of the asynchronous call, you pass the callbacks to the `begin_` method:

Copyright 2016, ZeroC, Inc.

Objective-C

```
e = [EmployeesPrx checkedCast:...]

[e begin_getName:99 response:getNameCB exception:failureCB];
```

You can also pass the Objective-C blocks directly to the call:

Objective-C

```
[e begin_getName:99
    response: ^(NSMutableString* name)
        {
            NSLog(@"Name is: %@", name);
        }
    exception: ^(ICEException* ex)
        {
            NSLog(@"Exception is: %@", [ex description]);
        }];
```

Ice enforces the following semantics at run time regarding which callbacks can be optionally specified with a `nil` value:

- You must supply an exception callback.
- You may omit the response callback for an operation that returns no data (that is, an operation with a void return type and no out-parameters).

Memory Management
The Ice run time creates an `NSAutoReleasePool` object before dispatching a completion callback. The pool is released once the dispatch is complete

Oneway Invocations in Objective-C

You can invoke operations via oneway proxies asynchronously, provided the operation has `void` return type, does not have any out-parameters, and does not raise user exceptions. If you call the `begin_` method on a oneway proxy for an operation that returns values or raises a user exception, the `begin_` method throws `NSException` with the `NSInvalidArgumentException` name.

The callback signatures look exactly as for a twoway invocation, but the response block is never called and may be `nil`.

Flow Control in Objective-C

Asynchronous method invocations never block the thread that calls the `begin_` method: the Ice run time checks to see whether it can write the request to the local transport. If it can, it does so immediately in the caller's thread. (In that case, `[ICEAsyncResult sentSynchronously]` returns true.) Alternatively, if the local transport does not have sufficient buffer space to accept the request, the Ice run time queues the request internally for later transmission in the background. (In that case, `[ICEAsyncResult sentSynchronously]` returns false.)

This creates a potential problem: if a client sends many asynchronous requests at the time the server is too busy to keep up with them, the requests pile up in the client-side run time until, eventually, the client runs out of memory.

Copyright 2016, ZeroC, Inc.

The API provides a way for you to implement flow control by counting the number of requests that are queued so, if that number exceeds some threshold, the client stops invoking more operations until some of the queued operations have drained out of the local transport.

You can supply a sent callback to be notified when the request was successfully sent:

Objective-C

```
void(^sentCB)(BOOL) = ^(BOOL sentSynchronously)
{
    ...
}
```

You inform the Ice run time that you want to be notified when a request has been passed to the local transport as usual:

Objective-C

```
[e begin_getName:99 response:getNameCB exception:failureCB sent:sentCB];
```

If the Ice run time can immediately pass the request to the local transport, it does so and invokes the sent callback from the thread that calls the `begin_` method. On the other hand, if the run time has to queue the request, it calls the sent callback from a different thread once it has written the request to the local transport. The boolean `sentSynchronously` parameter indicates whether the request was sent synchronously or was queued.

The sent callback allows you to limit the number of queued requests by counting the number of requests that are queued and decrementing the count when the Ice run time passes a request to the local transport.

Batch Requests in Objective-C

Applications that send batched requests can either flush a batch explicitly or allow the Ice run time to flush automatically. The proxy method `ice_flushBatchRequests` performs an immediate flush using the synchronous invocation model and may block the calling thread until the entire message can be sent. Ice also provides asynchronous versions of this method so you can flush batch requests asynchronously.

`begin_ice_flushBatchRequests` and `end_ice_flushBatchRequests` are proxy methods that flush any batch requests queued by that proxy.

In addition, similar methods are available on the communicator and the `Connection` object that is returned by `[ICEAsyncResult getConnection]`. These methods flush batch requests sent via the same communicator and via the same connection, respectively.

Concurrency in Objective-C

The Ice run time always invokes your callback methods from a separate thread, with one exception: it calls the sent callback from the thread calling the `begin_` method if the request could be sent synchronously. In the sent callback, you know which thread is calling the callback by looking at the `sentSynchronously` parameter.

See Also

- Request Contexts
- Batched Invocations

Copyright 2016, ZeroC, Inc.

slice2objc Command-Line Options

The Slice-to-Objective-C compiler, `slice2objc`, offers the following command-line options in addition to the standard options:

- `--include-dir DIR`
  Modifies `#import` directives in source files to prepend the path name of each header file with the directory *`DIR`*.
- `--dll-export` *`MACRO`*
  Use *`MACRO`* to control the export of symbols from dynamic shared libraries. This option allows you to export symbols from the generated code and place such generated code in a shared library. Compiling a Slice definition with:

  ```
  $ slice2objc --dll-export WIDGET_API x.ice
  ```

  adds the provided `MACRO` name (`WIDGET_API` in our example) to the declaration of interfaces and protocols that need to be exported.

  `The generated code also defines the provided MACRO as __attribute__((visibility ("default")))`.

  This option is useful when you create a shared library and compile your Objective-C code with `-fvisibility=hidden` to reduce the number of symbols exported.

See Also

- Using the Slice Compilers

Copyright 2016, ZeroC, Inc.

Example of a File System Client in Objective-C

This page presents a very simple client to access a server that implements the file system we developed in Slice for a Simple File System. The Objective-C code shown here hardly differs from the code you would write for an ordinary Objective-C program. This is one of the biggest advantages of using Ice: accessing a remote object is as easy as accessing an ordinary, local Objective-C object. This allows you to put your effort where you should, namely, into developing your application logic instead of having to struggle with arcane networking APIs. This is true for the server side as well, meaning that you can develop distributed applications easily and efficiently.

We now have seen enough of the client-side Objective-C mapping to develop a complete client to access our remote file system. For reference, here is the Slice definition once more:

Slice

```
["objc:prefix:FS"]
module Filesystem {
    exception GenericError {
        string reason;
    };

    interface Node {
        idempotent string name();
    };

    sequence<string> Lines;

    interface File extends Node {
        idempotent Lines read();
        idempotent void write(Lines text) throws GenericError;
    };

    sequence<Node*> NodeSeq;

    interface Directory extends Node {
        idempotent NodeSeq list();
    };
};
```

To exercise the file system, the client does a recursive listing of the file system, starting at the root directory. For each node in the file system, the client shows the name of the node and whether that node is a file or directory. If the node is a file, the client retrieves the contents of the file and prints them.

The body of the client code looks as follows:

Copyright 2016, ZeroC, Inc.

Objective-C

```
#import <Ice/Ice.h>
#import <Filesystem.h>

#import <Foundation/NSAutoreleasePool.h>
#import <stdio.h>

static void
listRecursive(id<FSDirectoryPrx> dir, int depth)
{
    // ...
}

int
main(int argc, char* argv[])
{
    NSAutoreleasePool * pool = [[NSAutoreleasePool alloc] init];
    int status = 1;
    id<ICECommunicator> communicator;
    @try
    {
        communicator = [ICEUtil createCommunicator:&argc argv:argv];

        // Create a proxy for the root directory
        //
        id<FSDirectoryPrx> rootDir = [FSDirectoryPrx checkedCast:
                [communicator stringToProxy: @"RootDir:default -p 10000"]];
        if (!rootDir)
            [NSException raise:@"invalid proxy" format:@"nil"];

        // Recursively list the contents of the root directory
        //
        printf("Contents of root directory:\n");
        listRecursive(rootDir, 0);

        status = 0;
    } @catch (NSException *ex) {
        NSLog(@"%@\n", [ex name]);
    }

    @try {
        [communicator destroy];
    } @catch (NSException* ex) {
        NSLog(@"%@\n", [ex name]);
    }

    [pool release];
    return status;
}
```

1. The code imports a few header files:

Copyright 2016, ZeroC, Inc.

- `Ice/Ice.h`: Always included in both client and server source files, provides definitions that are necessary for accessing the Ice run time.
- `Filesystem.h` The header that is generated by the Slice compiler from the Slice definitions in `Filesystem.ice`.
- `NSAutoreleasePool.h`: The client uses an autorelease pool to reclaim memory before it exits.
- `stdio.h`: The implementation of `listRecursive` prints to `stdout`.

2. The structure of the code in `main` follows what we saw in Hello World Application. After initializing the run time, the client creates a proxy to the root directory of the file system. For this example, we assume that the server runs on the local host and listens using the default protocol (TCP/IP) at port 10000. The object identity of the root directory is known to be `RootDir`.
3. The client down-casts the proxy to `DirectoryPrx` and passes that proxy to `listRecursive`, which prints the contents of the file system.

Most of the work happens in `listRecursive`:

Objective-C

```
// Print the specified number of tabs.

static void
printIndent(int depth)
{
    while (depth-- > 0)
        putchar('\t');
}

// Recursively print the contents of directory "dir" in tree
// fashion.  For files, show the contents of each file.
// The "depth" parameter is the current nesting level
// (for indentation).

static void
listRecursive(id<FSDirectoryPrx> dir, int depth)
{
    ++depth;
    FSNodeSeq *contents = [dir list];

    for (id<FSNodePrx> node in contents) {
        id<FSDirectoryPrx> dir = [FSDirectoryPrx checkedCast:node];
        id<FSFilePrx> file = [FSFilePrx uncheckedCast:node];
        printIndent(depth);
        printf("%s%s\n", [[node name] UTF8String],
(dir ? " (directory):" : " (file):"));
        if (dir) {
            listRecursive(dir, depth);
        } else {
            FSLines *text = [file read];
            for (NSString *line in text) {
                printIndent(depth);
                printf("\t%s\n", [line UTF8String]);
            }
        }
    }
}
```

The function is passed a proxy to a directory to list, and an indent level. (The indent level increments with each recursive call and allows the code to print the name of each node at an indent level that corresponds to the depth of the tree at that node.) `listRecursive` calls the `lis`

Copyright 2016, ZeroC, Inc.

`t` operation on the directory and iterates over the returned sequence of nodes:

1. The code does a `checkedCast` to narrow the `Node` proxy to a `Directory` proxy, as well as an `uncheckedCast` to narrow the `No de` proxy to a `File` proxy. Exactly one of those casts will succeed, so there is no need to call `checkedCast` twice: if the `Node` *is-a* `Directory`, the code uses the `id<FSDirectoryPrx>` returned by the `checkedCast`; if the `checkedCast` fails, we *know* that the `Node` *is-a* `File` and, therefore, an `uncheckedCast` is sufficient to get an `id<FSFilePrx>`.
   In general, if you know that a down-cast to a specific type will succeed, it is preferable to use an `uncheckedCast` instead of a `chec kedCast` because an `uncheckedCast` does not incur any network traffic.
2. The code prints the name of the file or directory and then, depending on which cast succeeded, prints `"(directory)"` or `"(file) "` following the name.
3. The code checks the type of the node:
   - If it is a directory, the code recurses, incrementing the indent level.
   - If it is a file, the code calls the `read` operation on the file to retrieve the file contents and then iterates over the returned sequence of lines, printing each line.

Assume that we have a small file system consisting of two files and a directory as follows:

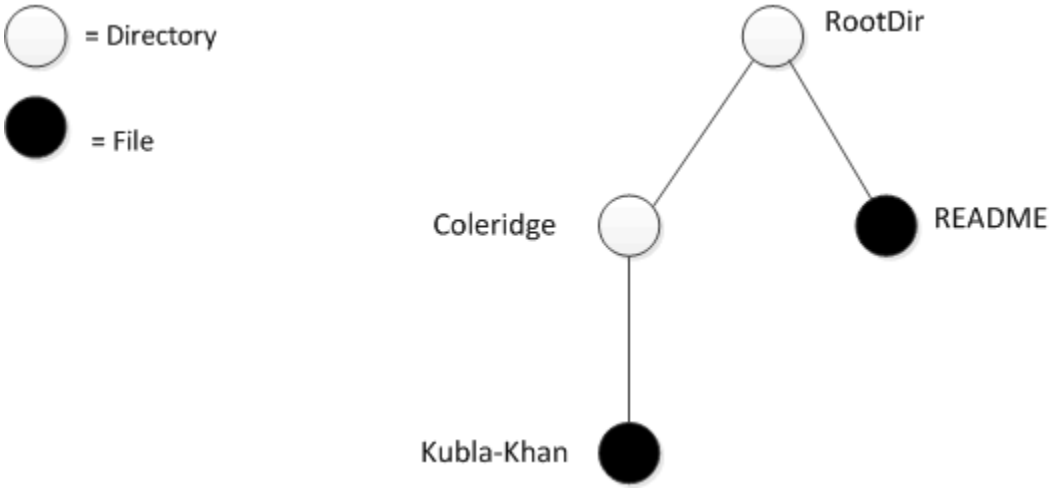

A small file system.

The output produced by the client for this file system is:

```
Contents of root directory:
        README (file):
                This file system contains a collection of poetry.
        Coleridge (directory):
                Kubla_Khan (file):
                        In Xanadu did Kubla Khan
                        A stately pleasure-dome decree:
                        Where Alph, the sacred river, ran
                        Through caverns measureless to man
                        Down to a sunless sea.
```

Note that, so far, our client (and server) are not very sophisticated:

- The protocol and address information are hard-wired into the code.
- The client makes more remote procedure calls than strictly necessary; with minor redesign of the Slice definitions, many of these calls can be avoided.

We will see how to address these shortcomings in our discussions of IceGrid and object life cycle.
See Also

- Hello World Application
- Slice for a Simple File System
- Example of a File System Server in Objective-C

Copyright 2016, ZeroC, Inc.

- Object Life Cycle
- IceGrid

Copyright 2016, ZeroC, Inc.

Server-Side Slice-to-Objective-C Mapping

The mapping for Slice data types to Objective-C is identical on the client side and server side, except for operation parameters, which map slightly differently for types that have mutable and immutable variants (strings, sequence, and dictionaries). This means that the mappings in the Client-Side Slice-to-Objective-C Mapping also apply to the server side. However, for the server side, there are a few additional things you need to know — specifically, how to:

- Initialize and finalize the server-side run time
- Implement servants
- Pass parameters and throw exceptions
- Create servants and register them with the Ice run time.

Although the examples we present are very simple, they accurately reflect the basics of writing an Ice server. Of course, for more sophisticated servers, you will be using additional APIs, for example, to improve performance or scalability. However, these APIs are all described in Slice, so, to use these APIs, you need not learn any Objective-C mapping rules beyond those described here.

Topics

- The Server-Side main Function in Objective-C
- Server-Side Objective-C Mapping for Interfaces
- Parameter Passing in Objective-C
- Raising Exceptions in Objective-C
- Object Incarnation in Objective-C
- Example of a File System Server in Objective-C

Copyright 2016, ZeroC, Inc.

The Server-Side main Function in Objective-C

This page discusses how to initialize and finalize the server-side run time.

On this page:

- Initializing and Finalizing the Server-Side Run Time
- Alternative Ways to Create a Communicator in Objective-C

Initializing and Finalizing the Server-Side Run Time

The main entry point to the Ice run time is represented by the local interface `ICECommunicator`. As for the client side, you must initialize the Ice run time by calling `createCommunicator` (a class method of the `ICEUtil` class) before you can do anything else in your server. `createCommunicator` returns an instance of type `id<ICECommunicator>`:

Objective-C

```
int
main(int argc, char* argv[])
{
    // ...
    id<ICECommunicator> communicator =
[ICEUtil createCommunicator:&argc argv:argv];
    // ...
}
```

`createCommunicator` accepts a *pointer* to `argc` as well as `argv`. The class method scans the argument vector for any command-line options that are relevant to the Ice run time; any such options are removed from the argument vector so, when `createCommunicator` returns, the only options and arguments remaining are those that concern your application. If anything goes wrong during initialization, `createCommunicator` throws an exception.

Before leaving your `main` function, you *must* call `Communicator::destroy`. The `destroy` operation is responsible for finalizing the Ice run time. In particular, `destroy` waits for any operation implementations that are still executing in the server to complete. In addition, `destroy` ensures that any outstanding threads are joined with and reclaims a number of operating system resources, such as file descriptors and memory. Never allow your `main` function to terminate without calling `destroy` first; doing so has undefined behavior.

The general shape of our server-side `main` function is therefore as follows:

Copyright 2016, ZeroC, Inc.

Objective-C

```
#import <Ice/Ice.h>

int
main(int argc, char* argv[])
{
    NSAutoreleasePool * pool = [[NSAutoreleasePool alloc] init];

    int status = 1;
    id<ICECommunicator> communicator = nil;
    @try {
        communicator = [ICEUtil createCommunicator:&argc argv:argv];

        // Server code here...

        status = 0;
    } @catch (NSException* ex) {
        NSLog(@"%@", ex);
    }

    @try {
        [communicator destroy];
    } @catch (NSException* ex) {
        NSLog(@"%@", ex);
    }

    [pool release];
    return status;
}
```

Note that the code places the call to `createCommunicator` into a `try` block and takes care to return the correct exit status to the operating system. Also note that the code creates and releases an autorelease pool. This ensures that memory will be released before the program terminates.

The `catch` handler for `NSException` ensures that the communicator is destroyed regardless of whether the program terminates normally or due to an exception.

You must *not* release the communicator that is returned by `createCommunicator`. As for any operation that returns a pointer, the Ice run time calls `autorelease` on the returned instance, so you do not have to release it yourself.

This is also the reason why `createCommunicator` is not called `initialize` (as it is for other language mappings) — `initialize` would suggest that the return value must be released because the method name begins with `init`.

Alternative Ways to Create a Communicator in Objective-C

`createCommunicator` is provided in several versions that accept different arguments. Here is the complete list:

- `(id<ICECommunicator>) createCommunicator: (int*)argc argv:(char*[])argv itData:(ICEInitializationData*)initData;`
  This is the designated initializer — the remaining versions of `createCommunicator` are implemented in terms of this initializer. As for the version we saw in the preceding section, this version accepts a pointer to `argc` as well as `argv` and removes Ice-related command-line options from the argument vector. The `initData` argument allows you to pass additional initialization information to the Ice run time (see below).

Copyright 2016, ZeroC, Inc.

- `+(id<ICECommunicator>) createCommunicator;`
  This is equivalent to calling:
  `[ICEUtil createCommunicator:nil argv:nil initData:nil];`
- `+(id<ICECommunicator>) createCommunicator: (int*)argc argv:(char*[])argv;`
  This is equivalent to calling
  `[ICEUtil createCommunicator:&argc argv:argv initData:nil];`
- `+(id<ICECommunicator>) createCommunicator: (ICEInitializationData*)initData`
  This is equivalent to calling
  `[ICEUtil createCommunicator:nil argv:nil initData:initData];`

The `initData` argument is of type `ICEInitializationData`. Even though it has no Slice definition, this class behaves as if it were a Slice structure with the following definition:

Slice

```
#include <Properties.ice>
#include <Logger.ice>

["objc:prefix:ICE"]
module Ice {
    dictionary<string, string> PrefixDict;

    local struct InitializationData {
        Ice::Properties properties;
        Ice::Logger logger;
    };
};
```

The `properties` member allows you to explicitly set property values for the communicator to be created. This is useful, for example, if you want to ensure that a particular property setting is always used by the communicator.

The `logger` member sets the logger that the Ice run time uses to log messages. If you do not set a logger (leaving the `logger` member as `nil`), the run time installs a default logger that calls `NSLog` to log messages.

See Also

- Communicator Initialization
- Properties and Configuration
- Logger Facility

Copyright 2016, ZeroC, Inc.

Server-Side Objective-C Mapping for Interfaces

The server-side mapping for interfaces provides an up-call API for the Ice run time: by implementing methods in a servant class, you provide the hook that gets the thread of control from the Ice server-side run time into your application code.

On this page:

- Skeleton Classes in Objective-C
- Servant Classes in Objective-C
  - Derived Servants in Objective-C
  - Delegate Servants in Objective-C
  - Memory Management

Skeleton Classes in Objective-C

On the client side, interfaces map to proxy protocols and classes. On the server side, interfaces map to *skeleton* protocols and classes. A skeleton is a class that has a method for each operation on the corresponding interface. For example, consider our Slice definition for the `Node` interface:

Slice

```
["objc:prefix:FS"]
module Filesystem {
    interface Node {
        idempotent string name();
    };
// ...
};
```

The Slice compiler generates the following definition for this interface:

Objective-C

```
@protocol FSNode <ICEObject>
-(NSString *) name:(ICECurrent *)current;
@end

@interface FSNode : ICEObject
// ...
@end
```

As you can see, the server-side API consists of a protocol and a class, known as the *skeleton protocol* and *skeleton class*. The methods of the skeleton class are internal to the mapping, so they do not concern us here. The skeleton protocol defines one method for each Slice operation. As for the client-side mapping, the method name is the same as the name of the corresponding Slice operation. If the Slice operation has parameters or a return value, these are reflected in the generated method, just as they are for the client-side mapping. In addition, each method has an additional trailing parameter of type `ICECurrent`. This parameter provides additional information about an invocation to your server-side code.

As for the client-side mapping, the generated code reflects the fact that all Slice interfaces and classes ultimately derive from `Ice::Object`. As you can see, the generated protocol incorporates the `ICEObject` protocol, and the generated class derives from the `ICEObject` class.

Servant Classes in Objective-C

The Objective-C mapping supports two different ways to implement servants. You can implement a servant by deriving from the skeleton

Copyright 2016, ZeroC, Inc.

class and implementing the methods for the Slice operations in your derived class. Alternatively, you can use a delegate servant, which need not derive from the skeleton class.

Derived Servants in Objective-C

To provide an implementation for an Ice object, you can create a servant class that derives from the corresponding skeleton class. For example, to create a servant for the `Node` interface, you could write:

Objective-C

```
@interface NodeI : FSNode <FSNode>
{
    @private
        NSString *myName;
}

+(id) nodei:(NSString *)name;
@end
```

By convention, servant classes have the name of their interface with an `I`-suffix, so the servant class for the `Node` interface is called `NodeI`. (This is a convention only: as far as the Ice run time is concerned, you can choose any name you prefer for your servant classes.)

Note that `NodeI` derives from `FSNode`, that is, it derives from its skeleton class. In addition, it adopts the `FSNode` protocol. Adopting the protocol is not strictly necessary; however, if you do write your servants this way, the compiler emits a warning if you forget to implement one or more Slice operations for the corresponding interface, so we suggest that you make it a habit to always have your servant class adopt its skeleton protocol.

As far as Ice is concerned, the `NodeI` class must implement the single `name` method that is defined by its skeleton protocol. That way, the Ice run time gets a servant that can respond to the operation that is defined by its Slice interface. You can add other methods and instance variables as you see fit to support your implementation. For example, in the preceding definition, we added a `myName` instance variable and property, a convenience constructor, and `dealloc`. Not surprisingly, the convenience constructor initializes the `myName` instance variable, the `name` method returns the value of that variable, and `dealloc` releases it:

Copyright 2016, ZeroC, Inc.

Objective-C

```
@implementation NodeI

+(id) nodei:(NSString *)name
{
    NodeI *instance = [[[NodeI alloc] init] autorelease];
    instance.myName = [[name copy] retain];
    return instance;
}

-(NSString *) name:(ICECurrent *)current
{
    return myName;
}

-(void) dealloc
{
    [myName release];
    [super dealloc];
}
@end
```

Delegate Servants in Objective-C

An alternate way to implement a servant is to use a delegate. `ICEObject` provides two constructors to do this:

Objective-C

```
@interface ICEObject : NSObject <ICEObject, NSCopying>
// ...
-(id) initWithDelegate:(id)delegate;
+(id) objectWithDelegate:(id)delegate;
@end
```

The `delegate` parameter specifies an object to which the servant will delegate operation invocations. That object need not derive from the skeleton class; the only requirement on the delegate is that it must have an implementation of the methods corresponding to the Slice operations that are called by clients. As for derived servants, we suggest that the delegate adopt the skeleton protocol, so the compiler will emit a warning if you forget to implement one or more Slice operations in the delegate.

The implementation of the Slice operations in a delegate servant is exactly the same as for a derived servant.

Delegate servants are useful if you need to derive your servant implementation from a base class in order to access some functionality. In that case, you cannot also derive the servant from the generated skeleton class. A delegate servant gets around Objective-C's single inheritance limitation and saves you having to write a servant class that forwards each operation invocation to the delegate.

Another use case are different interfaces that share their implementation. As an example, consider the following Slice definitions:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Intf1 {
    void op1();
};

interface Intf2 {
    void op2();
};
```

If `op1` and `op2` are substantially similar in their implementation and share common state, it can be convenient to implement the servants for `Intf1` and `Intf2` using a common delegate class:

Objective-C

```
@interface Intf1AndIntf2 : NSObject<EXIntf1, EXIntf2>
    +(id) intf1AndIntf2;
@end

@implementation Intf1AndIntf2
    +(id) intf1AndIntf2 { /*...*/ }
    -(void) op1:(ICECurrent*)current { /*...*/ }
    -(void) op2:(ICECurrent*)current { /*...*/ }
@end
```

See Instantiating an Objective-C Servant for an example of how to instantiate delegate servants.

Delegate servants do not permit you to override operations that are inherited from `ICEObject` (such as `ice_ping`). Therefore, if you want to override `ice_ping`, for example, to implement a default servant, you must use a derived servant.

Memory Management

The Ice run time creates an `NSAutoReleasePool` object before dispatching server-side invocations. The pool is released once the dispatch is complete.

See Also

- Slice for a Simple File System
- Objective-C Mapping for Interfaces
- Parameter Passing in Objective-C
- Object Incarnation in Objective-C
- The Current Object
- Default Servants

Copyright 2016, ZeroC, Inc.

Parameter Passing in Objective-C

This page shows how to implement parameters for Slice operations in Objective-C.

On this page:

- Implementing Parameters for Slice Operations in Objective-C
- Memory Management for Operations in Objective-C

Implementing Parameters for Slice Operations in Objective-C

For each parameter of a Slice operation, the Objective-C mapping generates a corresponding parameter for the method in the skeleton. In addition, every method has an additional, trailing parameter of type `ICECurrent`. For example, the `name` operation of the `Node` interface has no parameters, but the `name` method of the `Node` skeleton protocol has a single parameter of type `ICECurrent`. We will ignore this parameter for now.

Parameter passing on the server side follows the rules for the client side (with one exception):

- In-parameters and the return value are passed by value or by pointer, depending on the parameter type.
- Out-parameters are passed by pointer-to-pointer.

The exception to the client-side rules concerns types that come in mutable and immutable variants (strings, sequences, and dictionaries). For these, the server-side mapping passes the mutable variant where the client-side passes the immutable variant, and vice versa.

To illustrate the rules, consider the following interface that passes string parameters in all possible directions:

Slice

```
interface Intf {
    string op(string sin, out string sout);
};
```

The generated skeleton protocol for this interface looks as follows:

Objective-C

```
@protocol EXIntf <ICEObject>
-(NSString *) op:(NSMutableString *)sin
                 sout:(NSString **)sout
                 current:(ICECurrent *)current;
@end
```

As you can see, the in-parameter `sin` is of type `NSMutableString`, and the out parameter and return value are passed as `NSString` (the opposite of the client-side mapping). This means that in-parameters are passed to the servant as their mutable variant, and it is safe for you to modify such in-parameters. This is useful, for example, if a client passes a sequence to the operation, and the operation returns the sequence with a few minor changes. In that case, there is no need for the operation implementation to copy the sequence. Instead, you can simply modify the passed sequence as necessary and return the modified sequence to the client.

Here is an example implementation of the operation:

Copyright 2016, ZeroC, Inc.

Objective-C

```
-(NSString *) op:(NSMutableString *)sin
                 sout:(NSString **)sout
                 current:(ICECurrent *)current
{
    printf("%s\n", [sin UTF8String]); // In-params are initialized
    *sout = [sin appendString:@"appended"]; // Assign out-param
    return @"Done";                         // Return a string
}
```

Memory Management for Operations in Objective-C

To avoid leaking memory, you must be aware of how the Ice run time manages memory for operation implementations:

- In-parameters are passed to the servant already autoreleased.
- Out-parameters and return values must be returned by the servant as autoreleased values.

This follows the usual Objective-C convention: the allocator of a value is responsible for releasing it. This is what the Ice run time does for in-parameters, and what you are expected to do for out-parameters and return values. These rules also mean that it is safe to return an in-parameter as an out-parameter or return value. For example:

Objective-C

```
-(NSString *) op:(NSMutableString *)sin
                 sout:(NSString **)sout
                 current:(ICECurrent *)current
{
    *sout = sin; // Works fine.
    return sin;  // Works fine.
}
```

The Ice run time creates and releases a separate autorelease pool for each invocation. This means that the memory for parameters is reclaimed as soon as the run time has marshaled the operation results back to the client.
See Also

- Server-Side Objective-C Mapping for Interfaces
- Raising Exceptions in Objective-C
- The Current Object

Copyright 2016, ZeroC, Inc.

Raising Exceptions in Objective-C

To throw an exception from an operation implementation, you simply allocate the exception, initialize it, and throw it. For example:

Objective-C

```
-(void) write:(NSMutableArray *)text current:(ICECurrent *)current
{
    // Try to write the file contents here...
    // Assume we are out of space...
    if (error)
        @throw [FSGenericError genericError:@"file too large"];
}
```

As for out-parameters and return values, you must take care to throw an autoreleased exception.

If you throw an arbitrary Objective-C exception that does not derive from `ICEException`, the client receives an `UnknownException`. Similarly, if you throw an "impossible" user exception (a user exception that is not listed in the exception specification of the operation), the client receives an `UnknownUserException`.

If you throw a run-time exception, such as `MemoryLimitException`, the client receives an `UnknownLocalException`. For that reason, you should never throw system exceptions from operation implementations. If you do, all the client will see is an `UnknownLocalException`, which does not tell the client anything useful.

Three run-time exceptions are treated specially and not changed to `UnknownLocalException` when returned to the client: `ObjectNotExistException`, `OperationNotExistException`, and `FacetNotExistException`.

See Also

- Run-Time Exceptions
- Objective-C Mapping for Exceptions
- Server-Side Objective-C Mapping for Interfaces
- Parameter Passing in Objective-C

Copyright 2016, ZeroC, Inc.

Object Incarnation in Objective-C

Having created a servant class such as the rudimentary `NodeI` class, you can instantiate the class to create a concrete servant that can receive invocations from a client. However, merely instantiating a servant class is insufficient to incarnate an object. Specifically, to provide an implementation of an Ice object, you must follow these steps:

1. Instantiate a servant class.
2. Create an identity for the Ice object incarnated by the servant.
3. Inform the Ice run time of the existence of the servant.
4. Pass a proxy for the object to a client so the client can reach it.

On this page:

- Instantiating an Objective-C Servant
- Creating an Identity in Objective-C
- Activating an Objective-C Servant
- UUIDs as Identities in Objective-C
- Creating Proxies in Objective-C
  - Proxies and Servant Activation in Objective-C
  - Direct Proxy Creation in Objective-C

Instantiating an Objective-C Servant

Instantiating a servant means to allocate an instance on the heap:

Objective-C

```
NodeI *servant = [NodeI nodei:@"Fred"];
```

This code creates a new `NodeI` instance. For this example, we used the convenience constructor we saw earlier. Of course, you are not obliged to define such a constructor but, if you do not, you must explicitly call `release` or `autorelease` on the servant.

For a delegate servant, instantiation would look as follows:

Objective-C

```
Intf1AndIntf2 *delegate = [Intf1AndIntf2 intf1AndIntf2];
ICEObject *servant = [ICEObject objectWithDelegate:delegate];
```

Creating an Identity in Objective-C

Each Ice object requires an identity. That identity must be unique for all servants using the same object adapter.

The Ice object model assumes that all objects (regardless of their adapter) have a globally unique identity.

An Ice object identity is a structure with the following Slice definition:

Copyright 2016, ZeroC, Inc.

Slice

```
["objc:prefix:ICE"]
module Ice {
    struct Identity {
        string name;
        string category;
    };
    // ...
};
```

The full identity of an object is the combination of both the `name` and `category` fields of the `Identity` structure. For now, we will leave the `category` field as the empty string and simply use the `name` field. (The `category` field is most often used in conjunction with servant locators.)

To create an identity, we simply assign a key that identifies the servant to the `name` field of the `Identity` structure:

Objective-C

```
ICEIdentity ident = [ICEIdentity identity:"Fred" category:nil];
```

Activating an Objective-C Servant

Merely creating a servant instance does nothing: the Ice run time becomes aware of the existence of a servant only once you explicitly tell the object adapter about the servant. To activate a servant, you invoke the `add` operation on the object adapter. Assuming that we have access to the object adapter in the `adapter` variable, we can write:

Objective-C

```
[adapter add:servant identity:ident];
```

Note the two arguments to `add`: the servant and the object identity. Calling `add` on the object adapter adds the servant and the servant's identity to the adapter's servant map and links the proxy for an Ice object to the correct servant instance in the server's memory as follows:

1. The proxy for an Ice object, apart from addressing information, contains the identity of the Ice object. When a client invokes an operation, the object identity is sent with the request to the server.
2. The object adapter receives the request, retrieves the identity, and uses the identity as an index into the servant map.
3. If a servant with that identity is active, the object adapter retrieves the servant pointer from the servant map and dispatches the incoming request into the correct method on the servant.

Assuming that the object adapter is in the active state, client requests are dispatched to the servant as soon as you call `add`.

Putting the preceding points together, we can write a simple method that instantiates and activates one of our `NodeI` servants. For this example, we use a simple method on our servant called `activate` that activates a servant in an object adapter with the passed identity:

Copyright 2016, ZeroC, Inc.

Objective-C

```
-(void) activate:(id<ICEObjectAdapter>)a
                name:(NSString *)name
{
    ICEIdentity ident = [ICEIdentity identity:name category:nil];
    [a add:self identity:ident];
}
```

UUIDs as Identities in Objective-C

The Ice object model assumes that object identities are globally unique. One way of ensuring that uniqueness is to use UUIDs (Universally Unique Identifiers) as identities. The `ICEUtil` class contains a helper function to create such identities:

Objective-C

```
@interface ICEUtil : NSObject
+(id) generateUUID;
// ...
@end
```

When executed, this method returns a unique string such as `5029a22c-e333-4f87-86b1-cd5e0fcce509`. Each call to `generateUUID` creates a string that differs from all previous ones. You can use a UUID such as this to create object identities. For convenience, the object adapter has an operation `addWithUUID` that generates a UUID and adds a servant to the servant map in a single step:

Objective-C

```
-(id<ICEObjectPrx>) addWithUUID:(ICEObject*)servant
```

Note that the operation returns the proxy for the servant just activated.

Creating Proxies in Objective-C

Once we have activated a servant for an Ice object, the server can process incoming client requests for that object. However, clients can only access the object once they hold a proxy for the object. If a client knows the server's address details and the object identity, it can create a proxy from a string, as we saw in our first example. However, creation of proxies by the client in this manner is usually only done to allow the client access to initial objects for bootstrapping. Once the client has an initial proxy, it typically obtains further proxies by invoking operations.

The object adapter contains all the details that make up the information in a proxy: the addressing and protocol information, and the object identity. The Ice run time offers a number of ways to create proxies. Once created, you can pass a proxy to the client as the return value or as an out-parameter of an operation invocation.

Proxies and Servant Activation in Objective-C

The `add` and `addWithUUID` servant activation operations on the object adapter return a proxy for the corresponding Ice object, as we saw earlier. This means we can write:

Copyright 2016, ZeroC, Inc.

Objective-C

```
NodeI *servant = [NodeI nodei:name];
id<FSNodePrx> proxy = [FSNodePrx uncheckedCast: [adapter addWithUUID:servant]];

// Pass proxy to client...
```

Here, `addWithUUID` both activates the servant and returns a proxy for the Ice object incarnated by that servant in a single step.

Note that we need to use an `uncheckedCast` here because `addWithUUID` returns a proxy of type `id<ICEObjectPrx>`.

Direct Proxy Creation in Objective-C

The object adapter offers an operation to create a proxy for a given identity:

Slice

```
["objc:prefix:ICE"]
module Ice {
    local interface ObjectAdapter {
        Object* createProxy(Identity id);
        // ...
    };
};
```

Note that `createProxy` creates a proxy for a given identity whether a servant is activated with that identity or not. In other words, proxies have a life cycle that is quite independent from the life cycle of servants:

Objective-C

```
ICEIdentity *ident = [ICEIdentity identity];
ident.name = [ICEUtil generateUUID];
id<ICEObjectPrx> o = [adapter createProxy:ident];
```

This creates a proxy for an Ice object with the identity returned by `generateUUID`. Obviously, no servant yet exists for that object so, if we return the proxy to a client and the client invokes an operation on the proxy, the client will receive an `ObjectNotExistException`. (We examine these life cycle issues in more detail in Object Life Cycle.)

See Also

- Writing an Ice Application with Objective-C
- Server-Side Objective-C Mapping for Interfaces
- Object Adapter States
- Servant Locators
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

Example of a File System Server in Objective-C

This page presents the source code for a C++ server that implements our file system and communicates with the client we wrote earlier. The code here is fully functional, apart from the required interlocking for threads.

The server is remarkably free of code that relates to distribution: most of the server code is simply application logic that would be present just the same for a non-distributed version. Again, this is one of the major advantages of Ice: distribution concerns are kept away from application code so that you can concentrate on developing application logic instead of networking infrastructure.

The server code presented here is not quite correct as it stands: if two clients access the same file in parallel, each via a different thread, one thread may read the lines instance variable while another thread updates it. Obviously, if that happens, we may write or return garbage or, worse, crash the server. However, it is trivial to make the read and write operations thread-safe with a few lines of code. We discuss how to write thread-safe servant implementations in The Ice Threading Model.

On this page:

- Implementing a File System Server in Objective-C
- Server main Program in Objective-C
- Servant Class Definitions in Objective-C
- Servant Implementation in Objective-C
  - Implementing FileI in Objective-C
  - Implementing DirectoryI in Objective-C

Implementing a File System Server in Objective-C

We have now seen enough of the server-side Objective-C mapping to implement a server for our file system. (You may find it useful to review these Slice definitions before studying the source code.)

Our server is composed of three source files:

- `Server.m`
  This file contains the server main program.
- `FileI.m`
  This file contains the implementation for the `File` servants.
- `DirectoryI.m`
  This file contains the implementation for the `Directory` servants.

Server main Program in Objective-C

Our server main program, in the file `Server.m`, uses the structure we saw in an earlier example:

Objective-C

```
#import <Ice/Ice.h>
#import <FileI.h>
#import <DirectoryI.h>

#import <Foundation/NSAutoreleasePool.h>

int
main(int argc, char* argv[])
{
    NSAutoreleasePool * pool = [[NSAutoreleasePool alloc] init];

    int status = 1;
    id<ICECommunicator> communicator = nil;
    @try {
        communicator = [ICEUtil createCommunicator:&argc argv:argv];
```


Copyright 2016, ZeroC, Inc.

```
        id<ICEObjectAdapter> adapter =
[communicator createObjectAdapterWithEndpoints:
                                @"SimpleFilesystem"
                                endpoints:@"default -p 10000"];

        // Create the root directory (with name "/" and no parent)
        //
        DirectoryI *root = [DirectoryI directoryi:@"/" parent:nil];
        [root activate:adapter];

        // Create a file called "README" in the root directory
        //
        FileI *file = [FileI filei:@"README" parent:root];
        NSMutableArray *text = [NSMutableArray arrayWithObject:
            @"This file system contains a collection of poetry."];
        [file write:text current:nil];
        [file activate:adapter];

        // Create a directory called "Coleridge" in the root dir
        //
        DirectoryI *coleridge = [DirectoryI directoryi:@"Coleridge" parent:root];
        [coleridge activate:adapter];

        // Create a file called "Kubla_Khan"
        // in the Coleridge directory
        //
        file = [FileI filei:@"Kubla_Khan" parent:coleridge];
        text = [NSMutableArray arrayWithObjects:
                            @"In Xanadu did Kubla Khan",
                            @"A stately pleasure-dome decree:",
                            @"Where Alph, the sacred river, ran",
                            @"Through caverns measureless to man",
                            @"Down to a sunless sea.",
                            nil];
        [file write:text current:nil];
        [file activate:adapter];

        // All objects are created, allow client requests now
        //
        [adapter activate];

        // Wait until we are done
        //
        [communicator waitForShutdown];

        status = 0;
    } @catch (NSException* ex) {
        NSLog(@"%@", ex);
    }

    @try {
        [communicator destroy];
    } @catch (NSException* ex) {
        NSLog(@"%@", ex);
    }
```


Copyright 2016, ZeroC, Inc.

```
    [pool release];
```


Copyright 2016, ZeroC, Inc.

```
    return status;
}
```

There is quite a bit of code here, so let us examine each section in detail:

Objective-C

```
#import <Ice/Ice.h>
#import <FileI.h>
#import <DirectoryI.h>

#import <Foundation/NSAutoreleasePool.h>
```

The code includes the header `Ice/Ice.h`, which contains the definitions for the Ice run time, and the files `FileI.h` and `DirectoryI.h`, which contain the definitions of our servant implementations. Because we use an autorelease pool, we need to include `Foundation/NSAutoreleasePool.h` as well.

The next part of the source code is mostly boiler plate that we saw previously: we create an object adapter, and, towards the end, activate the object adapter and call `waitForShutdown`, which blocks the calling thread until you call `shutdown` or `destroy` on the communicator. (Ice does not make any demands on the main thread, so `waitForShutdown` simply blocks the calling thread; if you want to use the main thread for other purposes, you are free to do so.)

Copyright 2016, ZeroC, Inc.

Objective-C

```
int
main(int argc, char* argv[])
{
    NSAutoreleasePool * pool = [[NSAutoreleasePool alloc] init];

    int status = 1;
    id<ICECommunicator> communicator = nil;
    @try {
        communicator = [ICEUtil createCommunicator:&argc argv:argv];

        id<ICEObjectAdapter> adapter =
[communicator createObjectAdapterWithEndpoints:
                                @"SimpleFilesystem"
                                endpoints:@"default -p 10000"];

        // ...

        // All objects are created, allow client requests now
        //
        [adapter activate];

        // Wait until we are done
        //
        [communicator waitForShutdown];

        status = 0;
    } @catch (NSException* ex) {
        NSLog(@"%@", ex);
    }

    @try {
        [communicator destroy];
    } @catch (NSException* ex) {
        NSLog(@"%@", ex);
    }

    [pool release];
    return status;
}
```

The interesting part of the code follows the adapter creation: here, the server instantiates a few nodes for our file system to create the structure shown below:

Copyright 2016, ZeroC, Inc.

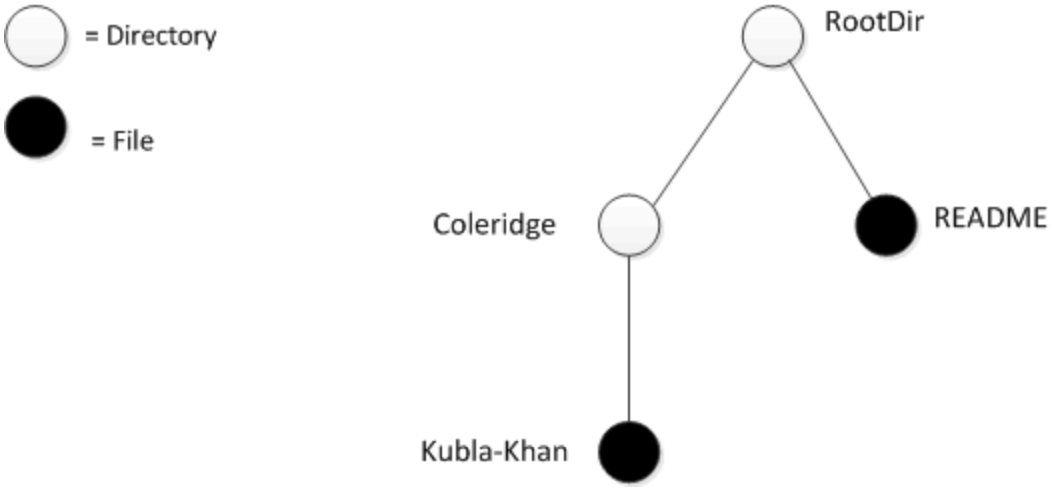

A small file system.

As we will see shortly, the servants for our directories and files are of type `DirectoryI` and `FileI`, respectively. The constructor for either type of servant accepts two parameters: the name of the directory or file to be created and the servant for the parent directory. (For the root directory, which has no parent, we pass a `nil` parent.) Thus, the statement

Objective-C

```
DirectoryI *root = [DirectoryI directoryi:@"/" parent:nil];
```

creates the root directory, with the name `"/"` and no parent directory.

Here is the code that establishes the structure in the above illustration shown:

Copyright 2016, ZeroC, Inc.

Objective-C

```
// Create the root directory (with name "/" and no parent)
//
DirectoryI *root = [DirectoryI directoryi:@"/" parent:nil];
[root activate:adapter];

// Create a file called "README" in the root directory
//
FileI *file = [FileI filei:@"README" parent:root];
NSMutableArray *text = [NSMutableArray arrayWithObject:
    @"This file system contains a collection of poetry."];
[file write:text current:nil];
[file activate:adapter];

// Create a directory called "Coleridge" in the root dir
//
DirectoryI *coleridge = [DirectoryI directoryi:@"Coleridge" parent:root];
[coleridge activate:adapter];

// Create a file called "Kubla_Khan"
// in the Coleridge directory
//
file = [FileI filei:@"Kubla_Khan" parent:coleridge];
text = [NSMutableArray arrayWithObjects:
                    @"In Xanadu did Kubla Khan",
                    @"A stately pleasure-dome decree:",
                    @"Where Alph, the sacred river, ran",
                    @"Through caverns measureless to man",
                    @"Down to a sunless sea.",
                    nil];
[file write:text current:nil];
[file activate:adapter];
```

We first create the root directory and a file `README` within the root directory. (Note that we pass the servant for the root directory as the parent pointer when we create the new node of type `FileI`.)

After creating each servant, the code calls `activate` on the servant. (We will see the definition of this member function shortly.) The `activate` member function adds the servant to the ASM.

The next step is to fill the file with text:

Copyright 2016, ZeroC, Inc.

Objective-C

```
// Create the root directory (with name "/" and no parent)
//
DirectoryI *root = [DirectoryI directoryi:@"/" parent:nil];
[root activate:adapter];

// Create a file called "README" in the root directory
//
FileI *file = [FileI filei:@"README" parent:root];
NSMutableArray *text = [NSMutableArray arrayWithObject:
    @"This file system contains a collection of poetry."];
[file write:text current:nil];
[file activate:adapter];

// Create a directory called "Coleridge" in the root dir
//
DirectoryI *coleridge = [DirectoryI directoryi:@"Coleridge" parent:root];
[coleridge activate:adapter];

// Create a file called "Kubla_Khan"
// in the Coleridge directory
//
file = [FileI filei:@"Kubla_Khan" parent:coleridge];
text = [NSMutableArray arrayWithObjects:
                    @"In Xanadu did Kubla Khan",
                    @"A stately pleasure-dome decree:",
                    @"Where Alph, the sacred river, ran",
                    @"Through caverns measureless to man",
                    @"Down to a sunless sea.",
                    nil];
[file write:text current:nil];
[file activate:adapter];
```

Recall that Slice sequences map to `NSArray` or `NSMutableArray`, depending on the parameter direction. Here, we instantiate that array and add a line of text to it.

Finally, we call the Slice `write` operation on our `FileI` servant by simply writing:

Objective-C

```
[file write:text current:nil];
```

This statement is interesting: the server code invokes an operation on one of its own servants. Because the call happens via the pointer to the servant (of type `FileI`) and *not* via a proxy (of type `id<FilePrx>`), the Ice run time does not know that this call is even taking place — such a direct call into a servant is not mediated by the Ice run time in any way and is dispatched as an ordinary Objective-C function call. The operation implementation in the servant expects a `current` object. In this case, we pass nil (which is fine because the operation implementation does not use it anyway).

In similar fashion, the remainder of the code creates a subdirectory called `Coleridge` and, within that directory, a file called `Kubla_Khan` to complete the structure in the above illustration.

Copyright 2016, ZeroC, Inc.

Servant Class Definitions in Objective-C

We must provide servants for the concrete interfaces in our Slice specification, that is, we must provide servants for the `File` and `Directory` interfaces in the Objective-C classes `FileI` and `DirectoryI`. This means that our servant classes look as follows:

Objective-C

```
#import <Filesystem.h>

@interface FileI : FSFile <FSFile>
// ...
@end

@interface DirectoryI : FSDirectory <FSDirectory>
// ...
@end
```

Each servant class derives from its skeleton class and adopts its skeleton protocol.

We now can think about how to implement our servants. One thing that is common to all nodes is that they have a name and a parent directory. As we saw earlier, we pass these details to a convenience constructor, which also takes care of calling `autorelease` on the new servant.

In addition, we will use UUIDs as the object identities for files and directories. This relieves us of the need to otherwise come up with a unique identity for each servant (such as path names, which would only complicate our implementation). Because the `list` operation returns proxies to nodes, and because each proxy carries the identity of the servant it denotes, this means that our servants must store their own identity, so we can create proxies to them when clients ask for them.

For `File` servants, we also need to store the contents of the file, leading to the following definition for the `FileI` class:

Copyright 2016, ZeroC, Inc.

Objective-C

```
#import <Filesystem.h>

@class DirectoryI;

@interface FileI : FSFile <FSFile>
{
    @private
        NSString *myName;
        DirectoryI *parent;
        ICEIdentity *ident;
        NSArray *lines;
}

@property(nonatomic, retain) NSString *myName;
@property(nonatomic, retain) DirectoryI *parent;
@property(nonatomic, retain) ICEIdentity *ident;
@property(nonatomic, retain) NSArray *lines;

+(id) filei:(NSString *)name parent:(DirectoryI *)parent;
-(void) write:(NSMutableArray *)text current:(ICECurrent *)current;
-(void) activate:(id<ICEObjectAdapter>)a;
@end
```

The instance variables store the name, parent node, identity, and the contents of the file. The `filei` convenience constructor instantiates the servant, remembers the name and parent directory, assigns a new identity, and calls `autorelease`.

Note that the only Slice operation we have defined here is the `write` method. This is necessary because, as we saw previously, the code in `Server.m` calls this method to initialize the files it creates.

For directories, the requirements are similar. They also need to store a name, parent directory, and object identity. Directories are also responsible for keeping track of the child nodes. We can store these nodes in an array of proxies. This leads to the following definition:

Copyright 2016, ZeroC, Inc.

Objective-C

```
#import <Filesystem.h>

@interface DirectoryI : FSDirectory <FSDirectory>
{
    @private
        NSString *myName;
        DirectoryI *parent;
        ICEIdentity *ident;
        NSMutableArray *contents;
}
@property(nonatomic, retain) NSString *myName;
@property(nonatomic, retain) DirectoryI *parent;
@property(nonatomic, retain) ICEIdentity *ident;
@property(nonatomic, retain) NSMutableArray *contents;

+(id) directoryi:(NSString *)name parent:(DirectoryI *)parent;
-(void) addChild:(id<FSNodePrx>)child;
-(void) activate:(id<ICEObjectAdapter>)a;
@end
```

Because the code in `Server.m` does not call any Slice operations on directory servants, we have not declared any of the corresponding methods. (We will see the purpose of the `addChild` method shortly.) As for files, the convenience constructor creates the servant, remembers the name and parent, and assigns an object identity, as well as calling `autorelease`.

Servant Implementation in Objective-C

Let us now turn to how to implement each of the methods for our servants.

Implementing `FileI` in Objective-C

The implementation of the `name`, `read`, and `write` operations for files is trivial, returning or updating the corresponding instance variable:

Objective-C

```
-(NSString *) name:(ICECurrent *)current
{
    return myName;
}

-(NSArray *) read:(ICECurrent *)current
{
    return lines;
}

-(void) write:(NSMutableArray *)text current:(ICECurrent *)current
{
    self.lines = text;
}
```


Copyright 2016, ZeroC, Inc.

Note that this constitutes the complete implementation of the Slice operations for files.

Here is the convenience constructor:

Objective-C

```
+(id) filei:(NSString *)name parent:(DirectoryI *)parent
{
    FileI *instance = [[[FileI alloc] init] autorelease];
    if(instance == nil)
    {
        return nil;
    }
    instance.myName = name;
    instance.parent = parent;
    instance.ident = [ICEIdentity identity:[ICEUtil generateUUID] category:nil];
    return instance;
}
```

After allocating and autoreleasing the instance, the constructor initializes the instance variables. The only interesting part of this code is how we create the identity for the servant. `generateUUID` is a class method of the `ICEUtil` class that returns a UUID. We assign this UUID to the `name` member of the identity.

We saw earlier that the server calls `activate` after it creates each servant. Here is the implementation of this method:

Objective-C

```
-(void) activate:(id<ICEObjectAdapter>)a
{
    id<FSNodePrx> thisNode = [FSNodePrx uncheckedCast:[a add:self identity:ident]];
    [parent addChild:thisNode];
}
```

This is how our code informs the Ice run time of the existence of a new servant. The call to `add` on the object adapter adds the servant and object identity to the adapter's servant map. In other words, this step creates the link between the object identity (which is embedded in proxies), and the actual Objective-C class instance that provides the behavior for the Slice operations.

`add` returns a proxy to the servant, of type `id<ICEObjectPrx>`. Because the `contents` instance variable of directory servants stores proxies of type `id<FSNodePrx>` (and `addChild` expects a proxy of that type), we down-cast the returned proxy to `id<FSNodePrx>`. In this case, because we know that the servant we just added to the adapter is indeed a servant that implements the operations on the Slice `Node` i nterface, we can use an `uncheckedCast`.

The call to `addChild` connects the new file to its parent directory.

Finally, we need a `dealloc` function so we do not leak the memory for the servant's instance variables:

Copyright 2016, ZeroC, Inc.

Objective-C

```
-(void) dealloc
{
    [myName release];
    [parent release];
    [ident release];
    [lines release];
    [super dealloc];
}
```

Implementing `DirectoryI` in Objective-C

The implementation of the Slice operations for directories is just as simple as for files:

Objective-C

```
-(NSString *) name:(ICECurrent *)current
{
    return myName;
}

-(NSArray *) list:(ICECurrent *)current
{
    return contents;
}
```

Because the `contents` instance variable stores the proxies for child nodes of the directory, the `list` operation simply returns that variable.

The convenience constructor looks much like the one for file servants:

Copyright 2016, ZeroC, Inc.

Objective-C

```
+(id) directoryi:(NSString *)name parent:(DirectoryI *)parent
{
    DirectoryI *instance = [[[DirectoryI alloc] init] autorelease];
    if(instance == nil)
    {
        return nil;
    }
    instance.myName = name;
    instance.parent = parent;
    instance.ident = [ICEIdentity
        identity:(parent ? [ICEUtil generateUUID] : @"RootDir")
        category:nil];
    instance.contents = [[NSMutableArray alloc] init];
    return instance;
}
```

The only noteworthy differences are that, for the root directory (which has no parent), the code uses `"RootDir"` as the identity. (As we saw earlier, the client knows that this is the identity of the root directory and uses it to create its proxy.)

The `addChild` method connects our nodes into a hierarchy by updating the `contents` instance variable. That way, each directory knows which nodes are contained in it:

Objective-C

```
-(void) addChild:(id<FSNodePrx>)child
{
     [contents addObject:child];
}
```

Finally, the `activate` and `dealloc` methods are very much like the corresponding methods for files:

Copyright 2016, ZeroC, Inc.

Objective-C

```
-(void) activate:(id<ICEObjectAdapter>)a
{
    id<FSNodePrx> thisNode = [FSNodePrx uncheckedCast:[a add:self identity:ident]];
    [parent addChild:thisNode];
}

-(void) dealloc
{
    [myName release];
    [parent release];
    [ident release];
    [contents release];
    [super dealloc];
}
```

See Also

- Slice for a Simple File System
- Objective-C Mapping for Sequences
- Example of a File System Client in Objective-C
- The Server-Side main Function in Objective-C
- The Ice Threading Model

Copyright 2016, ZeroC, Inc.

PHP Mapping

Ice currently provides a client-side mapping for PHP, but not a server-side mapping.

Topics

- Client-Side Slice-to-PHP Mapping

Copyright 2016, ZeroC, Inc.

Client-Side Slice-to-PHP Mapping

The client-side Slice-to-PHP mapping defines how Slice data types are translated to PHP types, and how clients invoke operations, pass parameters, and handle errors. Much of the PHP mapping is intuitive. For example, Slice sequences map to PHP arrays, so there is essentially nothing new you have to learn in order to use Slice sequences in PHP.

Much of what appears in this chapter is reference material. We suggest that you skim the material on the initial reading and refer back to specific sections as needed. However, we recommend that you read at least the mappings for exceptions, interfaces, and operations in detail because these sections cover how to call operations from a client, pass parameters, and handle exceptions.

In order to use the PHP mapping, you should need no more than the Slice definition of your application and knowledge of the PHP mapping rules. In particular, looking through the generated code in order to discern how to use the PHP mapping is likely to be inefficient, due to the amount of detail. Of course, occasionally, you may want to refer to the generated code to confirm a detail of the mapping, but we recommend that you otherwise use the material presented here to see how to write your client-side code.

The `Ice` Module
All of the APIs for the Ice run time are nested in the `Ice` module, to avoid clashes with definitions for other libraries or applications. Some of the contents of the `Ice` module are generated from Slice definitions; other parts of the `Ice` module provide special-purpose definitions that do not have a corresponding Slice definition. We will incrementally cover the contents of the `Ice` m odule throughout the remainder of the manual.

A PHP application can load the Ice run time using the `require` statement:

```
require 'Ice.php';
```

If the statement executes without error, the Ice run time is loaded and available for use. You can determine the version of the Ice run time you have just loaded by calling the `stringVersion` function:

```
$icever = Ice_stringVersion();
```

Using the namespace mapping, you can refer to a global Ice function such as `stringVersion` either by its flattened name (as shown above) or by its namespace equivalent:

```
$icever = \Ice\stringVersion();
```

Topics

- PHP Mapping for Identifiers
- PHP Mapping for Modules
- PHP Mapping for Built-In Types
- PHP Mapping for Enumerations
- PHP Mapping for Structures
- PHP Mapping for Sequences
- PHP Mapping for Dictionaries
- PHP Mapping for Constants
- PHP Mapping for Exceptions
- PHP Mapping for Interfaces
- PHP Mapping for Operations
- PHP Mapping for Classes
- slice2php Command-Line Options
- Application Notes for PHP
- Using Slice Checksums in PHP
- Example of a File System Client in PHP

Copyright 2016, ZeroC, Inc.

PHP Mapping for Identifiers

A Slice identifier maps to an identical PHP identifier. For example, the Slice identifier `Clock` becomes the PHP identifier `Clock`. There is one exception to this rule: if a Slice identifier is the same as a PHP keyword or is an identifier reserved by the Ice run time (such as `checked Cast`), the corresponding PHP identifier is prefixed with an underscore. For example, the Slice identifier `while` is mapped as `_while`.

You should try to avoid such identifiers as much as possible.

A single Slice identifier often results in several PHP identifiers. For example, for a Slice interface named `Foo`, the generated PHP code uses the identifiers `Foo` and `FooPrx` (among others). If the interface has the name `while`, the generated identifiers are `_while` and `whilePrx` (*not* `_whilePrx`), that is, the underscore prefix is applied only to those generated identifiers that actually require it.

See Also

- Lexical Rules
- PHP Mapping for Modules
- PHP Mapping for Built-In Types
- PHP Mapping for Enumerations
- PHP Mapping for Structures
- PHP Mapping for Sequences
- PHP Mapping for Dictionaries
- PHP Mapping for Constants
- PHP Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

PHP Mapping for Modules

By default, identifiers defined within a Slice module are mapped to a flattened symbol that uses underscores as module separators. Consider the following Slice definition:

Slice

```
module M {
    module N {
        enum Color { red, green, blue };
    };
};
```

The Slice identifier `Color` maps to `M_N_Color` by default because PHP releases prior to version 5.3 lacked language support for namespaces. If you prefer to use namespaces instead, you can enable an alternate mapping in which Slice modules map to PHP namespaces with the same name as the Slice module. This mapping preserves the nesting of the Slice definitions. Using the namespace mapping, the Slice identifier `Color` maps to `\M\N\Color`.

Be aware that using underscores in your Slice definitions can lead to name collisions in the flattened mapping. Consider the following example:

Slice

```
module M {
    module N {
        enum Color { red, green, blue };
    };
};

module M_N {
    interface Color { };
};
```

Although these definitions are syntactically correct, they both map to the flattened PHP symbol `M_N_Color`.
See Also

- Modules
- PHP Mapping for Identifiers
- PHP Mapping for Built-In Types
- PHP Mapping for Enumerations
- PHP Mapping for Structures
- PHP Mapping for Sequences
- PHP Mapping for Dictionaries
- PHP Mapping for Constants
- PHP Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

PHP Mapping for Built-In Types

On this page:

- Mapping of Slice Built-In Types to PHP Types
- String Mapping in PHP

Mapping of Slice Built-In Types to PHP Types

PHP has a limited set of primitive types: `boolean`, `integer`, `double`, and `string`. The Slice built-in types are mapped to PHP types as shown in the table below:

Slice	Ruby
`bool`	`true` or `false`
`byte`	`integer`
`short`	`integer`
`int`	`integer`
`long`	`integer`
`float`	`double`
`double`	`double`
`string`	`string`

PHP's `integer` type may not accommodate the range of values supported by Slice's `long` type, therefore `long` values that are outside this range are mapped as strings. Scripts must be prepared to receive an integer or string from any operation that returns a `long` value.

String Mapping in PHP

String values returned as the result of a Slice operation (including return values, out parameters, and data members) contain UTF-8 encoded strings unless the program has installed a string converter, in which case string values use the converter's native encoding instead.

As string input values for a remote Slice operation, Ice accepts `null` in addition to `string` objects; each occurrence of `null` is marshaled as an empty string. Ice assumes that all `string` objects contain valid UTF-8 encoded strings unless the program has installed a string converter, in which case Ice assumes that `string` objects use the native encoding expected by the converter.

See Also

- Basic Types
- PHP Mapping for Identifiers
- PHP Mapping for Modules
- PHP Mapping for Enumerations
- PHP Mapping for Structures
- PHP Mapping for Sequences
- PHP Mapping for Dictionaries
- PHP Mapping for Constants
- PHP Mapping for Exceptions
- C++ Strings and Character Encoding

Copyright 2016, ZeroC, Inc.

PHP Mapping for Enumerations

PHP does not have an enumerated type, so a Slice enumeration is mapped to a PHP class: the name of the Slice enumeration becomes the name of the PHP class; for each enumerator, the class contains a constant with the same name as the enumerator. For example:

Slice

```
enum Fruit { Apple, Pear, Orange };
```

The generated PHP class looks as follows:

PHP

```
class Fruit
{
    const Apple = 0;
    const Pear = 1;
    const Orange = 2;
}
```

Suppose we modify the Slice definition to include a custom enumerator value:

Slice

```
enum Fruit { Apple, Pear = 3, Orange };
```

The generated PHP class changes accordingly:

PHP

```
class Fruit
{
    const Apple = 0;
    const Pear = 3;
    const Orange = 4;
}
```

Since enumerated values are mapped to integer constants, application code is not required to use the generated constants. When an enumerated value enters the Ice run time, Ice validates that the given integer is a valid value for the enumeration. However, to minimize the potential for defects in your code, we recommend using the generated constants instead of literal integers.

See Also

- Enumerations
- PHP Mapping for Identifiers
- PHP Mapping for Modules
- PHP Mapping for Built-In Types

Copyright 2016, ZeroC, Inc.

- PHP Mapping for Structures
- PHP Mapping for Sequences
- PHP Mapping for Dictionaries
- PHP Mapping for Constants
- PHP Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

PHP Mapping for Structures

A Slice structure maps to a PHP class containing a public variable for each member of the structure. For example, here is our `Employee` structure once more:

Slice

```
struct Employee {
    long number;
    string firstName;
    string lastName;
};
```

The PHP mapping generates the following definition for this structure:

PHP

```
class Employee
{
    public function __construct($number=0, $firstName='', $lastName='');

    public function __toString();

    public $number;
    public $firstName;
    public $lastName;
}
```

The class provides a constructor whose arguments correspond to the data members. This allows you to instantiate and initialize the class in a single statement (instead of having to first instantiate the class and then assign to its members). Each argument provides a default value appropriate for the member's type:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	Null
`dictionary`	Null
`class/interface`	Null

You can also declare different default values for members of primitive and enumerated types.

The mapping also includes a definition for the `__toString` magic method, which returns a string representation of the structure.

See Also

Copyright 2016, ZeroC, Inc.

- Structures
- PHP Mapping for Identifiers
- PHP Mapping for Modules
- PHP Mapping for Built-In Types
- PHP Mapping for Enumerations
- PHP Mapping for Sequences
- PHP Mapping for Dictionaries
- PHP Mapping for Constants
- PHP Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

PHP Mapping for Sequences

A Slice sequence maps to a native PHP indexed array. The first element of the Slice sequence is contained at index 0 (zero) of the PHP array, followed by the remaining elements in ascending index order.

Consider this example:

Slice

```
sequence<Fruit> FruitPlatter;
```

You can create an instance of `FruitPlatter` as shown below:

PHP

```
// Make a small platter with one Apple and one Orange
//
$platter = array(Fruit::Apple, Fruit::Orange);
```

The Ice run time validates the elements of an array to ensure that they are compatible with the declared type and raises `InvalidArgument Exception` if an incompatible type is encountered.
See Also

- Sequences
- PHP Mapping for Identifiers
- PHP Mapping for Modules
- PHP Mapping for Built-In Types
- PHP Mapping for Enumerations
- PHP Mapping for Structures
- PHP Mapping for Dictionaries
- PHP Mapping for Constants
- PHP Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

PHP Mapping for Dictionaries

A Slice dictionary maps to a native PHP associative array. The PHP mapping does not currently support all Slice dictionary types, however, because native PHP associative arrays support only integers and strings as keys.

A Slice dictionary whose key type is an enumeration or one of the primitive types `boolean`, `byte`, `short`, `int`, or `long` is mapped as an associative array with an integer key.

Boolean values are treated as integers, with false equivalent to 0 (zero) and true equivalent to 1 (one).

A Slice dictionary with a `string` key type is mapped as an associative array with a string key. All other key types cause a warning to be generated.

Here is the definition of our `EmployeeMap`:

Slice

```
dictionary<long, Employee> EmployeeMap;
```

You can create an instance of this dictionary as shown below:

PHP

```
$e1 = new Employee;
$e1->number = 42;
$e1->firstName = "Stan";
$e1->lastName = "Lipmann";

$e2 = new Employee;
$e2->number = 77;
$e2->firstName = "Herb";
$e2->lastName = "Sutter";

$em = array($e1->number => $e1, $e2->number => $e2);
```

The Ice run time validates the elements of a dictionary to ensure that they are compatible with the declared type; `InvalidArgumentException` exception is raised if an incompatible type is encountered.

See Also

- Dictionaries
- PHP Mapping for Identifiers
- PHP Mapping for Modules
- PHP Mapping for Built-In Types
- PHP Mapping for Enumerations
- PHP Mapping for Structures
- PHP Mapping for Sequences
- PHP Mapping for Constants
- PHP Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

PHP Mapping for Constants

A Slice constant maps to a PHP constant. Consider the following definitions:

Slice

```
module M {
    const bool      AppendByDefault = true;
    const byte      LowerNibble = 0x0f;
    const string    Advice = "Don't Panic!";
    const short     TheAnswer = 42;
    const double    PI = 3.1416;

    enum Fruit { Apple, Pear, Orange };
    const Fruit     FavoriteFruit = Pear;
};
```

The mapping for these constants is shown below:

PHP

```
define('M_AppendByDefault', true);
define('M_LowerNibble', 15);
define('M_Advice', "Don't Panic!");
define('M_TheAnswer', 42);
define('M_PI', 3.1416);
define('M_FavoriteFruit', M_Fruit::Pear);
```

Slice string literals that contain non-ASCII characters or universal character names are mapped to PHP string literals with these characters replaced by their UTF-8 encoding as octal escapes. For example:

Slice

```
module M {
   const string Egg = "œuf";
   const string Heart = "c\u0153ur";
   const string Banana = "\U0001F34C";
};
```

is mapped to:

PHP

```
define('M_Egg', "\305\223uf");
define('M_Heart', "c\305\223ur");
define('M_Banana', "\360\237\215\214");
```

An application refers to a constant using its flattened name:

Copyright 2016, ZeroC, Inc.

PHP

```
$ans = M_TheAnswer;
```

Using the namespace mapping, Slice constants are mapped to PHP constants in the enclosing namespace:

PHP

```
$ans = \M\TheAnswer;
```

See Also

- Constants and Literals
- PHP Mapping for Identifiers
- PHP Mapping for Modules
- PHP Mapping for Built-In Types
- PHP Mapping for Enumerations
- PHP Mapping for Structures
- PHP Mapping for Sequences
- PHP Mapping for Dictionaries
- PHP Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

PHP Mapping for Exceptions

On this page:

- Inheritance Hierarchy for Exceptions in PHP
- PHP Mapping for User Exceptions
  - Optional Data Members
- PHP Mapping for Run-Time Exceptions

Inheritance Hierarchy for Exceptions in PHP

The mapping for exceptions is based on the inheritance hierarchy shown below:

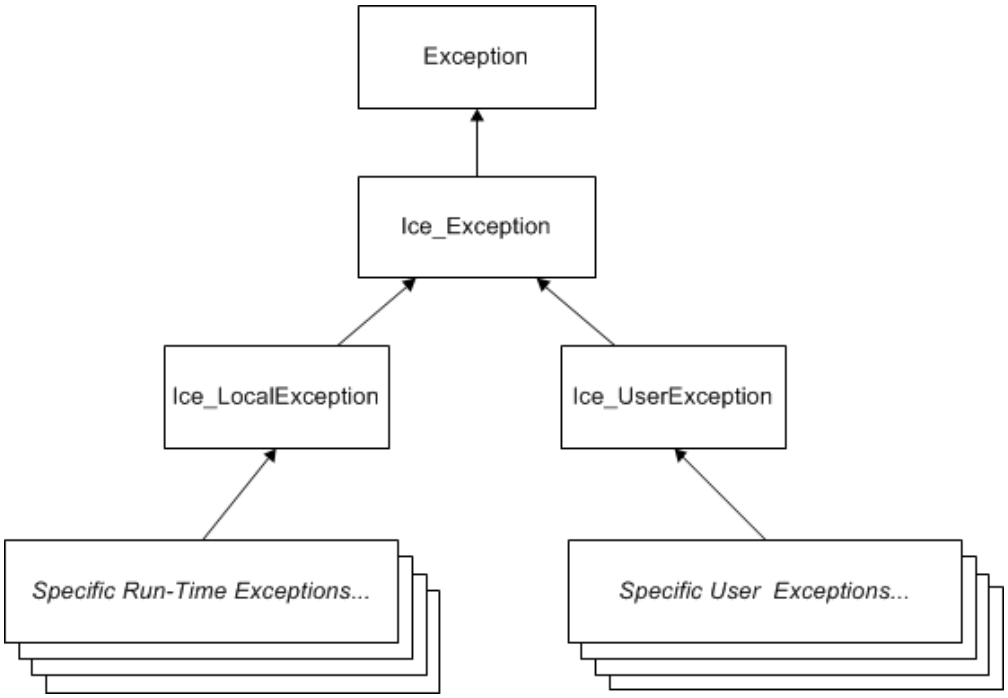

Inheritance structure for Ice exceptions.

The ancestor of all exceptions is `Exception`, from which `Ice_Exception` is derived. `Ice_LocalException` and `Ice_UserException` are derived from `Ice_Exception` and form the base for all run-time and user exceptions.

PHP Mapping for User Exceptions

Here is a fragment of the Slice definition for our world time server once more:

Slice

```
exception GenericError {
    string reason;
};
exception BadTimeVal extends GenericError {};
exception BadZoneName extends GenericError {};
```

These exception definitions map to the abbreviated PHP class definitions shown below:

Copyright 2016, ZeroC, Inc.

PHP

```
class GenericError extends Ice_UserException
{
    public function __construct($reason='');
    public function ice_name();
    public function __toString();

    public $reason;
}

class BadTimeVal extends GenericError
{
    public function __construct($reason='');
    public function ice_name();
    public function __toString();
}

class BadZoneName extends GenericError
{
    public function __construct($reason='');
    public function ice_name();
    public function __toString();
}
```

Each Slice exception is mapped to a PHP class with the same name. The inheritance structure of the Slice exceptions is preserved for the generated classes, so `BadTimeVal` and `BadZoneName` inherit from `GenericError`.

Each exception member corresponds to an instance variable of the instance, which the constructor initializes to a default value appropriate for its type:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	Null
`dictionary`	Null
`class/interface`	Null

You can also declare different default values for members of primitive and enumerated types. For derived exceptions, the constructor has one parameter for each of the base exception's data members, plus one parameter for each of the derived exception's data members, in base-to-derived order. As an example, although `BadTimeVal` and `BadZoneName` do not declare data members, their constructors still accept a value for the inherited data member `reason` in order to pass it to the constructor of the base exception `GenericError`.

Each exception also defines the `ice_name` method to return the exception's type name, as well as the `__toString` magic method to return a stringified representation of the exception and its members.

All user exceptions are derived from the base class `Ice_UserException`. This allows you to catch all user exceptions generically by

Copyright 2016, ZeroC, Inc.

installing a handler for `Ice_UserException`. Similarly, you can catch all Ice run-time exceptions with a handler for `Ice_LocalException`, and you can catch all Ice exceptions with a handler for `Ice_Exception`.

Optional Data Members

Optional data members use the same mapping as required data members, but an optional data member can also be set to the marker value `Ice_Unset` to indicate that the member is unset. A well-behaved program must compare an optional data member to `Ice_Unset` before using the member's value:

PHP

```
try {
    ...
} catch($ex) {
    if($ex->optionalMember == Ice_Unset)
        echo "optionalMember is unset\n";
    else
        echo "optionalMember = " . $ex->optionalMember . "\n";
}
```

The `Ice_Unset` marker value has different semantics than `null`. Since `null` is a legal value for certain Slice types, the Ice run time requires a separate marker value so that it can determine whether an optional value is set. An optional value set to `null` is considered to be set. If you need to distinguish between an unset value and a value set to `null`, you can do so as follows:

PHP

```
try {
    ...
} catch($ex) {
    if($ex->optionalMember == Ice_Unset)
        echo "optionalMember is unset\n";
    else if($ex->optionalMember == null)
        echo "optionalMember is None\n";
    else
        echo "optionalMember = " . $ex->optionalMember . "\n";
}
```

PHP Mapping for Run-Time Exceptions

The Ice run time throws run-time exceptions for a number of pre-defined error conditions. All run-time exceptions directly or indirectly derive from `Ice_LocalException` (which, in turn, derives from `Ice_Exception`).

By catching exceptions at the appropriate point in the inheritance hierarchy, you can handle exceptions according to the category of error they indicate:

- `Ice_LocalException`
  This is the root of the inheritance tree for run-time exceptions.
- `Ice_UserException`
  This is the root of the inheritance tree for user exceptions.
- `Ice_TimeoutException`
  This is the base exception for both operation-invocation and connection-establishment timeouts.

Copyright 2016, ZeroC, Inc.

- `Ice_ConnectTimeoutException`
  This exception is raised when the initial attempt to establish a connection to a server times out.

For example, `Ice_ConnectTimeoutException` can be handled as `Ice_ConnectTimeoutException`, `Ice_TimeoutException`, `Ice_LocalException`, or `Ice_Exception`.

You will probably have little need to catch run-time exceptions as their most-derived type and instead catch them as `Ice_LocalException`; the fine-grained error handling offered by the remainder of the hierarchy is of interest mainly in the implementation of the Ice run time. Exceptions to this rule are the exceptions related to facet and object life cycles, which you may want to catch explicitly. These exceptions are `Ice_FacetNotExistException` and `Ice_ObjectNotExistException`, respectively.

See Also

- User Exceptions
- Run-Time Exceptions
- PHP Mapping for Identifiers
- PHP Mapping for Modules
- PHP Mapping for Built-In Types
- PHP Mapping for Enumerations
- PHP Mapping for Structures
- PHP Mapping for Sequences
- PHP Mapping for Dictionaries
- PHP Mapping for Constants
- Optional Data Members
- Versioning
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

PHP Mapping for Interfaces

The mapping of Slice interfaces revolves around the idea that, to invoke a remote operation, you call a member function on a local class instance that is a proxy for the remote object. This makes the mapping easy and intuitive to use because making a remote procedure call is no different from making a local procedure call (apart from error semantics).

On this page:

- Proxy Objects in PHP
- Ice_ObjectPrx Class in PHP
- Proxy Helper Classes in PHP
- Casting Proxies in PHP
    - Proxy Backward Compatibility in PHP
- Using Proxy Methods in PHP
- Object Identity and Proxy Comparison in PHP

Proxy Objects in PHP

Slice interfaces are implemented by instances of the `Ice_ObjectPrx` class. In the client's address space, an instance of `ObjectPrx` is the local ambassador for a remote instance of an interface in a server and is known as a *proxy instance*. All the details about the server-side object, such as its address, what protocol to use, and its object identity are encapsulated in that instance.

The PHP mapping for proxies differs from that of other Ice language mappings in that the `ObjectPrx` class is used to implement *all* Slice interfaces. The primary motivation for this design is minimizing the amount of code that is generated for each interface. As a result, a proxy object returned by the communicator operations `stringToProxy` and `propertyToProxy` is *untyped*, meaning it is not associated with a user-defined Slice interface. Once you narrow the proxy to a particular interface, you can use that proxy to invoke your Slice operations.

Proxy instances are always created on behalf of the client by the Ice run time, so client code never has any need to instantiate a proxy directly.

A value of `null` denotes the null proxy. The null proxy is a dedicated value that indicates that a proxy points "nowhere" (denotes no object).

For each operation in the interface, the proxy object supports a method of the same name. Each operation accepts an optional trailing parameter representing the operation context. This parameter is an associative string array for use by the Ice run time to store information about how to deliver a request. You normally do not need to use it. (We examine the context parameter in detail in Request Contexts. The parameter is also used by IceStorm.)

`Ice_ObjectPrx` Class in PHP

In the PHP language mapping, all proxies are instances of `Ice_ObjectPrx`. This class provides a number of methods:

PHP

```
class Ice_ObjectPrx
{
    function ice_getIdentity();
    function ice_isA($id);
    function ice_ids();
    function ice_id();
    function ice_ping();
    # ...
}
```

The methods behave as follows:

- **`ice_getIdentity`**
  This method returns the identity of the object denoted by the proxy. The identity of an Ice object has the following Slice type:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    struct Identity {
        string name;
        string category;
    };
};
```

To see whether two proxies denote the same object, first obtain the identity for each object and then compare the identities:

PHP

```
$proxy1 = ...
$proxy2 = ...
$id1 = $proxy1->ice_getIdentity();
$id2 = $proxy2->ice_getIdentity();

if($id1 == $id2)
    // proxy1 and proxy2 denote the same object
else
    // proxy1 and proxy2 denote different objects
```

- **ice_isA**
  The ice_isA method determines whether the object denoted by the proxy supports a specific interface. The argument to ice_isA is a type ID. For example, to see whether a proxy of type ObjectPrx denotes a Printer object, we can write:

  PHP

  ```
  $proxy = ...
  if($proxy != null && $proxy->ice_isA("::Printer"))
      // proxy denotes a Printer object
  else
      // proxy denotes some other type of object
  ```

  Note that we are testing whether the proxy is null before attempting to invoke the ice_isA method. This avoids getting a run-time error if the proxy is null.

- **ice_ids**
  The ice_ids method returns an array of strings representing all of the type IDs that the object denoted by the proxy supports.

- **ice_id**
  The ice_id method returns the type ID of the object denoted by the proxy. Note that the type returned is the type of the actual object, which may be more derived than the static type of the proxy. For example, if we have a proxy of type BasePrx, with a static type ID of ::Base, the return value of ice_id might be "::Base", or it might be something more derived, such as "::Derived".

- **ice_ping**
  The ice_ping method provides a basic reachability test for the object. If the object can physically be contacted (that is, the object exists and its server is running and reachable), the call completes normally; otherwise, it throws an exception that indicates why the

Copyright 2016, ZeroC, Inc.

object could not be reached, such as `ObjectNotExistException` or `ConnectTimeoutException`.

The `ObjectPrx` class also defines an operator for comparing two proxies for equality. Note that all aspects of proxies are compared by this operation, such as the communication endpoints for the proxy. This means that, in general, if two proxies compare unequal, that does *not* imply that they denote different objects. For example, if two proxies denote the same Ice object via different transport endpoints, `==` returns `false` even though the proxies denote the same object.

The `ice_isA`, `ice_ids`, `ice_id`, and `ice_ping` methods are remote operations and therefore support an additional overloading that accepts a request context. Also note that there are other methods in `ObjectPrx`, not shown here. These methods provide different ways to dispatch a call and also provide access to an object's facets.

Proxy Helper Classes in PHP

The PHP mapping for a proxy generates a helper class with several static methods. For example, the following class is generated for the Slice interface named `Simple`:

PHP

```
class SimplePrxHelper
{
    public static function checkedCast($proxy, $facetOrCtx=null, $ctx=null);

    public static function uncheckedCast($proxy, $facet=null);

    public static function ice_staticId();
}
```

The `checkedCast` and `uncheckedCast` methods are described in the following section.

The `ice_staticId` method returns the type ID string corresponding to the interface. As an example, for the Slice interface `Simple` in module `M`, the string returned by `ice_staticId` is `"::M::Simple"`.

Casting Proxies in PHP

As shown above, the proxy's helper class includes two static methods that support down-casting:

PHP

```
class SimplePrxHelper
{
    public static function checkedCast($proxy, $facetOrCtx=null, $ctx=null);
    public static function uncheckedCast($proxy, $facet=null);

    // ...
}
```

The method names `checkedCast` and `uncheckedCast` are reserved for use in proxies. If a Slice interface defines an operation with either of those names, the mapping escapes the name in the generated proxy by prepending an underscore. For example, an interface that defines an operation named `checkedCast` is mapped to a proxy with a method named `_checkedCast`.

For `checkedCast`, if the passed proxy is for an object of type `Simple`, or a proxy for an object with a type derived from `Simple`, the cast returns a proxy narrowed to that type; otherwise, if the passed proxy denotes an object of a different type (or if the passed proxy is `null`), the cast returns `null`.

The arguments are described below:

Copyright 2016, ZeroC, Inc.

- `$proxy`
  The proxy to be narrowed.
- `$facetOrCtx`
  This optional argument can be either a string representing a desired facet, or an associative string array representing a context.
- `$ctx`
  If `$facetOrCtx` contains a facet name, use this argument to supply an associative string array representing a context.
- `$facet`
  Specifies the name of the desired facet.

Given a proxy of any type, you can use a `checkedCast` to determine whether the corresponding object supports a given type, for example:

PHP

```
$obj = ...      // Get a proxy from somewhere...

$simple = SimplePrxHelper::checkedCast($obj);
if($simple != null)
    // Object supports the Simple interface...
else
    // Object is not of type Simple...
```

Note that a `checkedCast` contacts the server. This is necessary because only the server implementation has definite knowledge of the type of an object. As a result, a `checkedCast` may throw a `ConnectTimeoutException` or an `ObjectNotExistException`.

In contrast, an `uncheckedCast` does not contact the server and unconditionally returns a proxy of the requested type. However, if you do use an `uncheckedCast`, you must be certain that the proxy really does support the type you are casting to; otherwise, if you get it wrong, you will most likely get a run-time exception when you invoke an operation on the proxy. The most likely error for such a type mismatch is `OperationNotExistException`. However, other exceptions, such as a marshaling exception are possible as well. And, if the object happens to have an operation with the correct name, but different parameter types, no exception may be reported at all and you simply end up sending the invocation to an object of the wrong type; that object may do rather nonsensical things. To illustrate this, consider the following two interfaces:

Slice

```
interface Process {
    void launch(int stackSize, int dataSize);
};

// ...

interface Rocket {
    void launch(float xCoord, float yCoord);
};
```

Suppose you expect to receive a proxy for a `Process` object and use an `uncheckedCast` to down-cast the proxy:

Copyright 2016, ZeroC, Inc.

PHP

```
$obj = ...                                          // Get proxy...
$process = ProcessPrxHelper::uncheckedCast($obj); // No worries...
$process->launch(40, 60);                           // Oops...
```

If the proxy you received actually denotes a `Rocket` object, the error will go undetected by the Ice run time: because `int` and `float` have the same size and because the Ice protocol does not tag data with its type on the wire, the implementation of `Rocket::launch` will simply misinterpret the passed integers as floating-point numbers.

In fairness, this example is somewhat contrived. For such a mistake to go unnoticed at run time, both objects must have an operation with the same name and, in addition, the run-time arguments passed to the operation must have a total marshaled size that matches the number of bytes that are expected by the unmarshaling code on the server side. In practice, this is extremely rare and an incorrect `uncheckedCast` typically results in a run-time exception.

Proxy Backward Compatibility in PHP

Prior releases of the PHP language mapping provided two proxy methods for narrowing a proxy:

PHP

```
class Ice_ObjectPrx
{
    function ice_checkedCast($type, $facetOrCtx=null, $ctx=null);
    function ice_uncheckedCast($type, $facet=null);
    # ...
}
```

For example, a proxy can be narrowed as follows:

PHP

```
$proxy = $proxy->ice_checkedCast("::Demo::Hello");
```

Embedding such type ID strings in your application is a potential source of defects because the strings are not validated until run time. Although these methods are still supported for the sake of backward compatibility, we recommend using the static methods that are generated in the helper class corresponding to each interface, as shown below:

PHP

```
$proxy = Demo_HelloPrxHelper::checkedCast($proxy);
```

Not only are these static methods consistent with the APIs of other Ice language mappings, they also avoid the need to hard-code type ID strings in your application.

Copyright 2016, ZeroC, Inc.

Using Proxy Methods in PHP

The base proxy class `ObjectPrx` supports a variety of methods for customizing a proxy. Since proxies are immutable, each of these "factory methods" returns a copy of the original proxy that contains the desired modification. For example, you can obtain a proxy configured with a ten second timeout as shown below:

PHP

```
$proxy = $communicator->stringToProxy(...);
$proxy = $proxy->ice_timeout(10000);
```

A factory method returns a new proxy object if the requested modification differs from the current proxy, otherwise it returns the current proxy. With few exceptions, factory methods return a proxy of the same type as the current proxy, therefore it is generally not necessary to repeat a down-cast after using a factory method. The example below demonstrates these semantics:

PHP

```
$base = $communicator->stringToProxy(...);
$hello = Demo_HelloPrxHelper::checkedCast($base);
$hello = $hello->ice_timeout(10000); // Type is not discarded
$hello->sayHello();
```

The only exceptions are the factory methods `ice_facet` and `ice_identity`. Calls to either of these methods may produce a proxy for an object of an unrelated type, therefore they return an untyped proxy that you must subsequently down-cast to an appropriate type.

Object Identity and Proxy Comparison in PHP

Proxy objects support comparison using the comparison operators == and `!=`. Note that proxy comparison uses *all* of the information in a proxy for the comparison. This means that not only the object identity must match for a comparison to succeed, but other details inside the proxy, such as the protocol and endpoint information, must be the same. In other words, comparison tests for *proxy* identity, *not* object identity. A common mistake is to write code along the following lines:

PHP

```
$p1 = ...        // Get a proxy...
$p2 = ...        // Get another proxy...

if($p1 != $p2)
    // p1 and p2 denote different objects       // WRONG!
else
    // p1 and p2 denote the same object         // Correct
```

Even though `p1` and `p2` differ, they may denote the same Ice object. This can happen because, for example, both `p1` and `p2` embed the same object identity, but each uses a different protocol to contact the target object. Similarly, the protocols may be the same, but denote different endpoints (because a single Ice object can be contacted via several different transport endpoints). In other words, if two proxies compare equal, we know that the two proxies denote the same object (because they are identical in all respects); however, if two proxies compare unequal, we know absolutely nothing: the proxies may or may not denote the same object.

To compare the object identities of two proxies, you can use helper functions in the `Ice` module:

Copyright 2016, ZeroC, Inc.

PHP

```
function Ice_proxyIdentityCompare($lhs, $rhs);
function Ice_proxyIdentityAndFacetCompare($lhs, $rhs);
```

`proxyIdentityCompare` allows you to correctly compare proxies for identity:

PHP

```
$p1 = ...        // Get a proxy...
$p2 = ...        // Get another proxy...

if(Ice_proxyIdentityCompare($p1, $p2) != 0)
    // p1 and p2 denote different objects       // Correct
else
    // p1 and p2 denote the same object         // Correct
```

The function returns 0 if the identities are equal, `-1` if `p1` is less than `p2`, and 1 if `p1` is greater than `p2`. (The comparison uses `name` as the major sort key and `category` as the minor sort key.)

The `proxyIdentityAndFacetCompare` function behaves similarly, but compares both the identity and the facet name.

See Also

- Interfaces, Operations, and Exceptions
- Proxies for Ice Objects
- Type IDs
- PHP Mapping for Operations
- Request Contexts
- Versioning
- IceStorm

Copyright 2016, ZeroC, Inc.

PHP Mapping for Operations

On this page:

- Basic PHP Mapping for Operations
- Normal and idempotent Operations in PHP
- Passing Parameters in PHP
  - In-Parameters in PHP
  - Out-Parameters in PHP
  - Parameter Type Mismatches in PHP
  - Null Parameters in PHP
  - Optional Parameters in PHP
- Exception Handling in PHP

Basic PHP Mapping for Operations

As we saw in the PHP mapping for interfaces, for each operation on an interface, a proxy object narrowed to that type supports a corresponding method with the same name. To invoke an operation, you call it via the proxy. For example, here is part of the definitions for our file system:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };
    // ...
};
```

The `name` operation returns a value of type `string`. Given a proxy to an object of type `Node`, the client can invoke the operation as follows:

PHP

```
$node = ...             // Initialize proxy
$name = $node->name();  // Get name via RPC
```

Normal and `idempotent` Operations in PHP

You can add an `idempotent` qualifier to a Slice operation. As far as the signature for the corresponding proxy method is concerned, `idempotent` has no effect.

Passing Parameters in PHP

In-Parameters in PHP

The PHP mapping for `in` parameters guarantees that the value of a parameter will not be changed by the invocation.

Here is an interface with operations that pass parameters of various types from client to server:

Copyright 2016, ZeroC, Inc.

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ClientToServer {
    void op1(int i, float f, bool b, string s);
    void op2(NumberAndString ns, StringSeq ss, StringTable st);
    void op3(ClientToServer* proxy);
};
```

A proxy object narrowed to the `ClientToServer` interface supports the following methods:

PHP

```
function op1($i, $f, $b, $s, $_ctx=null);
function op2($ns, $ss, $st, $_ctx=null);
function op3($proxy, $_ctx=null);
```

Given a proxy to a `ClientToServer` interface, the client code can pass parameters as in the following example:

Copyright 2016, ZeroC, Inc.

PHP

```
$p = ...                                  // Get proxy...

$p->op1(42, 3.14, true, "Hello world!"); // Pass simple literals

$i = 42;
$f = 3.14;
$b = true;
$s = "Hello world!";
$p->op1($i, $f, $b, $s);                  // Pass simple variables

$ns = new NumberAndString;
$ns->x = 42;
$ns->str = "The Answer";
$ss = array("Hello world!");
$st = array();
$st[0] = $ns;
$p->op2($ns, $ss, $st);                   // Pass complex variables

$p->op3($p);                              // Pass proxy
```

Out-Parameters in PHP

Out parameters are passed by reference. Here is the same Slice definition we saw earlier, but this time with all parameters being passed in the out direction:

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ServerToClient {
    int op1(out float f, out bool b, out string s);
    void op2(out NumberAndString ns,
             out StringSeq ss,
             out StringTable st);
    void op3(out ServerToClient* proxy);
};
```

The PHP mapping looks the same as it did for the in-parameters version:

Copyright 2016, ZeroC, Inc.

PHP

```
function op1($i, $f, $b, $s, $_ctx=null);
function op2($ns, $ss, $st, $_ctx=null);
function op3($proxy, $_ctx=null);
```

Given a proxy to a `ServerToClient` interface, the client code can receive the results as in the following example:

PHP

```
$p = ...                 // Get proxy...
$p->op1($i, $f, $b, $s);
$p->op2($ns, $ss, $st);
$p->op3($stcp);
```

Note that it is not necessary to use the reference operator (&) before each argument because the Ice run time forces each `out` parameter to have reference semantics.

Parameter Type Mismatches in PHP

The Ice run time performs validation on the arguments to a proxy invocation and reports any type mismatches as `InvalidArgumentException`.

Null Parameters in PHP

Some Slice types naturally have "empty" or "not there" semantics. Specifically, sequences, dictionaries, and strings all can be `null`, but the corresponding Slice types do not have the concept of a null value. To make life with these types easier, whenever you pass `null` as a parameter or return value of type sequence, dictionary, or string, the Ice run time automatically sends an empty sequence, dictionary, or string to the receiver.

This behavior is useful as a convenience feature: especially for deeply-nested data types, members that are sequences, dictionaries, or strings automatically arrive as an empty value at the receiving end. This saves you having to explicitly initialize, for example, every string element in a large sequence before sending the sequence in order to avoid a run-time error. Note that using null parameters in this way does *not* create null semantics for Slice sequences, dictionaries, or strings. As far as the object model is concerned, these do not exist (only *empty* sequences, dictionaries, and strings do). For example, it makes no difference to the receiver whether you send a string as `null` or as an empty string: either way, the receiver sees an empty string.

Optional Parameters in PHP

Optional parameters use the same mapping as required parameters. The only difference is that `Ice_Unset` can be passed as the value of an optional parameter or return value. Consider the following operation:

Slice

```
optional(1) int execute(optional(2) string params, out optional(3) float value);
```

A client can invoke this operation as shown below:

Copyright 2016, ZeroC, Inc.

PHP

```
$i = $proxy->execute("--file log.txt", $v);
$i = $proxy->execute(Ice_Unset, $v);

if($v != Ice_Unset)
    echo "value = " . $v . "\n";
```

A well-behaved program must always compare an optional parameter to `Ice_Unset` prior to using its value. Keep in mind that the `Ice_Unset` marker value has different semantics than `null`. Since `null` is a legal value for certain Slice types, the Ice run time requires a separate marker value so that it can determine whether an optional parameter is set. An optional parameter set to `null` is considered to be set. If you need to distinguish between an unset parameter and a parameter set to `null`, you can do so as follows:

PHP

```
if($optionalParam == Ice_Unset)
    echo "optionalParam is unset\n";
else if($optionalParam == null)
    echo "optionalParam is null\n";
else
    echo "optionalParam = " . $optionalParam . "\n";
```

Exception Handling in PHP

Any operation invocation may throw a run-time exception and, if the operation has an exception specification, may also throw user exceptions. Suppose we have the following simple interface:

Slice

```
exception Tantrum {
    string reason;
};

interface Child {
    void askToCleanUp() throws Tantrum;
};
```

Slice exceptions are thrown as PHP exceptions, so you can simply enclose one or more operation invocations in a `try-catch` block:

Copyright 2016, ZeroC, Inc.

PHP

```
$child = ...        // Get child proxy...

try
{
    $child->askToCleanUp();
}
catch(Tantrum $t)
{
    echo "The child says: " . $t->reason . "\n";
}
```

Typically, you will catch only a few exceptions of specific interest around an operation invocation; other exceptions, such as unexpected run-time errors, will usually be handled by exception handlers higher in the hierarchy. For example:

PHP

```
function run()
{
    $child = ...          // Get child proxy...
    try
    {
        $child->askToCleanUp();
    }
    catch(Tantrum $t)
    {
        echo "The child says: " . $t->reason . "\n";
        $child->scold();  // Recover from error...
    }
    $child->praise();     // Give positive feedback...
}

try
{
    // ...
    run();
    // ...
}
catch(Ice_Exception $ex)
{
    echo $ex->__toString() . "\n";
}
```

This code handles a specific exception of local interest at the point of call and deals with other exceptions generically. (This is also the strategy we used for our first simple application in Hello World Application.)

See Also

- Operations
- Hello World Application

Copyright 2016, ZeroC, Inc.

- Slice for a Simple File System
- PHP Mapping for Interfaces
- PHP Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

PHP Mapping for Classes

On this page:

- Basic PHP Mapping for Classes
- Inheritance from Object in PHP
- Class Data Members in PHP
- Class Constructors in PHP
- Class Operations in PHP
- Class Factories in PHP

Basic PHP Mapping for Classes

A Slice class maps to a Ruby class with the same name. For each Slice data member, the generated class contains a member variable, just as for structures and exceptions. Consider the following class definition:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The PHP mapping generates the following code for this definition:

Copyright 2016, ZeroC, Inc.

PHP

```
abstract class TimeOfDay extends Ice_ObjectImpl
{
    public function __construct($hour=0, $minute=0, $second=0)
    {
        $this->hour = $hour;
        $this->minute = $minute;
        $this->second = $second;
    }

    abstract public function format();

    public static function ice_staticId()
    {
        return '::TimeOfDay';
    }

    public function __toString()
    {
        // ...
    }

    public $hour;
    public $minute;
    public $second;
}
```

There are a number of things to note about the generated code:

1. The generated class `TimeOfDay` inherits from `Ice_ObjectImpl`. This reflects the semantics of Slice classes in that all classes implicitly inherit from `Object`, which is the ultimate ancestor of all classes. Note that `Object` is *not* the same as `Ice_ObjectPrx`. In other words, you *cannot* pass a class where a proxy is expected and vice versa.
2. The constructor initializes an instance variable for each Slice data member.
3. The class includes an abstract function declaration corresponding to the Slice operation `format`.
4. The class defines the class method `ice_staticId`.

There is quite a bit to discuss here, so we will look at each item in turn.

Inheritance from `Object` in PHP

Like interfaces, classes implicitly inherit from a common base class, `Ice_Object`. However, classes inherit from `Ice_Object` instead of `Ice_ObjectPrx`, therefore you cannot pass a class where a proxy is expected (and vice versa) because the base types for classes and proxies are not compatible.

`Ice_Object` contains a number of member functions:

Copyright 2016, ZeroC, Inc.

PHP

```
interface Ice_Object
{
    public function ice_isA($id);

    public function ice_ping();

    public function ice_ids();

    public function ice_id();

    public function ice_preMarshal();

    public function ice_postUnmarshal();
}
```

The member functions of `Ice_Object` behave as follows:

- `ice_isA`
  This method returns `true` if the object supports the given type ID, and `false` otherwise.
- `ice_ping`
  As for interfaces, `ice_ping` provides a basic reachability test for the object.
- `ice_ids`
  This method returns a string sequence representing all of the type IDs supported by this object, including `::Ice::Object`.
- `ice_id`
  This method returns the actual run-time type ID of the object. If you call `ice_id` through a reference to a base instance, the returned type ID is the actual (possibly more derived) type ID of the instance.
- `ice_preMarshal`
  If the object supports this method, the Ice run time invokes it just prior to marshaling the object's state, providing the opportunity for the object to validate its declared data members.
- `ice_postUnmarshal`
  If the object supports this method, the Ice run time invokes it after unmarshaling the object's state. An object typically defines this method when it needs to perform additional initialization using the values of its declared data members.

All Slice classes derive from `Ice_Object` via the `Ice_ObjectImpl` abstract base class, which provides default implementations of the `Ice_Object` methods.

Class Data Members in PHP

By default, data members of classes are mapped exactly as for structures and exceptions: for each data member in the Slice definition, the generated class contains a corresponding member variable.

Optional data members use the same mapping as required data members, but an optional data member can also be set to the marker value `Ice_Unset` to indicate that the member is unset. A well-behaved program must compare an optional data member to `Ice_Unset` before using the member's value:

Copyright 2016, ZeroC, Inc.

PHP

```
$v = ...;
if($v->optionalMember == Ice_Unset)
    echo "optionalMember is unset\n";
else
    echo "optionalMember = " . $v->optionalMember . "\n";
```

The `Ice_Unset` marker value has different semantics than `null`. Since `null` is a legal value for certain Slice types, the Ice run time requires a separate marker value so that it can determine whether an optional value is set. An optional value set to `null` is considered to be set. If you need to distinguish between an unset value and a value set to `null`, you can do so as follows:

PHP

```
$v = ...;
if($v->optionalMember == Ice_Unset)
    echo "optionalMember is unset\n";
else if($v->optionalMember == null)
    echo "optionalMember is None\n";
else
    echo "optionalMember = " . $v->optionalMember . "\n";
}
```

If you wish to restrict access to a data member, you can modify its visibility using the `protected` metadata directive. The presence of this directive causes the Slice compiler to generate the data member with protected visibility. As a result, the member can be accessed only by the class itself or by one of its subclasses. For example, the `TimeOfDay` class shown below has the `protected` metadata directive applied to each of its data members:

Slice

```
class TimeOfDay {
    ["protected"] short hour;   // 0 - 23
    ["protected"] short minute; // 0 - 59
    ["protected"] short second; // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Slice compiler produces the following generated code for this definition:

Copyright 2016, ZeroC, Inc.

PHP

```
abstract class TimeOfDay extends Ice_ObjectImpl
{
    public function __construct($hour=0, $minute=0, $second=0)
    {
        $this->hour = $hour;
        $this->minute = $minute;
        $this->second = $second;
    }

    abstract public function format();

    public static function ice_staticId()
    {
        return '::TimeOfDay';
    }

    public function __toString()
    {
        // ...
    }

    protected $hour;
    protected $minute;
    protected $second;
}
```

For a class in which all of the data members are protected, the metadata directive can be applied to the class itself rather than to each member individually. For example, we can rewrite the `TimeOfDay` class as follows:

Slice

```
["protected"] class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

Class Constructors in PHP

Classes have a constructor that assigns to each data member a default value appropriate for its type:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration

Copyright 2016, ZeroC, Inc.

`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	Null
`dictionary`	Null
`class/interface`	Null

You can also declare different default values for data members of primitive and enumerated types.

For derived classes, the constructor has one parameter for each of the base class's data members, plus one parameter for each of the derived class's data members, in base-to-derived order.

Pass the marker value `Ice_Unset` as the value of any optional data members that you wish to be unset.

Class Operations in PHP

Operations of classes are mapped to abstract member functions in the generated class. This means that, if a class contains operations (such as the `format` operation of our `TimeOfDay` class), you must provide an implementation of the operation in a class that is derived from the generated class. For example:

PHP

```
class TimeOfDayI extends TimeOfDay
{
    public function format()
    {
        return strftime("%X");
    }
}
```

Class Factories in PHP

Having created a class such as `TimeOfDayI`, we have an implementation and we can instantiate the `TimeOfDayI` class, but we cannot receive it as the return value or as an out-parameter from an operation invocation. To see why, consider the following simple interface:

Slice

```
interface Time {
    TimeOfDay get();
};
```

When a client invokes the `get` operation, the Ice run time must instantiate and return an instance of the `TimeOfDay` class. However, `TimeOfDay` is an abstract class that cannot be instantiated. Unless we tell it, the Ice run time cannot magically know that we have created a `TimeOfDayI` class that implements the abstract `format` operation of the `TimeOfDay` abstract class. In other words, we must provide the Ice run time with a factory that knows that the `TimeOfDay` abstract class has a `TimeOfDayI` concrete implementation. The `Ice::Communicator` interface provides us with the necessary operations:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local interface ObjectFactory {
        Object create(string type);
        void destroy();
    };

    local interface Communicator {
        void addObjectFactory(ObjectFactory factory, string id);
        ObjectFactory findObjectFactory(string id);
        // ...
    };
};
```

To supply the Ice run time with a factory for our `TimeOfDayI` class, we must implement the `ObjectFactory` interface:

PHP

```
class ObjectFactory implements Ice_ObjectFactory {
    public function create($type) {
        if ($type == TimeOfDay::ice_staticId()) {
            return new TimeOfDayI;
        }
        assert(false);
        return null;
    }

    public function destroy() {
        // Nothing to do
    }
}
```

The object factory's `create` method is called by the Ice run time when it needs to instantiate a `TimeOfDay` class. The factory's `destroy` m ethod is called by the Ice run time when its communicator is destroyed.

The `create` method is passed the type ID of the class to instantiate. For our `TimeOfDay` class, the type ID is `"::TimeOfDay"`. Our implementation of `create` checks the type ID: if it matches, the method instantiates and returns a `TimeOfDayI` object. For other type IDs, the method asserts because it does not know how to instantiate other types of objects.

Note that we used the `ice_staticId` method to obtain the type ID rather than embedding a literal string. Using a literal type ID string in your code is discouraged because it can lead to errors that are only detected at run time. For example, if a Slice class or one of its enclosing modules is renamed and the literal string is not changed accordingly, a receiver will fail to unmarshal the object and the Ice run time will raise `NoObjectFactoryException`. By using `ice_staticId` instead, we avoid any risk of a misspelled or obsolete type ID, and we can discover at compile time if a Slice class or module has been renamed.

Given a factory implementation, such as our `ObjectFactory`, we must inform the Ice run time of the existence of the factory:

Copyright 2016, ZeroC, Inc.

PHP

```
$communicator = ...;
$communicator->addObjectFactory(new ObjectFactory, TimeOfDay::ice_staticId());
```

Now, whenever the Ice run time needs to instantiate a class with the type ID `"::TimeOfDay"`, it calls the `create` method of the registered `ObjectFactory` instance.

The `destroy` operation of the object factory is invoked by the Ice run time when the communicator is destroyed. This gives you a chance to clean up any resources that may be used by your factory. Do not call `destroy` on the factory while it is registered with the communicator — if you do, the Ice run time has no idea that this has happened and, depending on what your `destroy` implementation is doing, may cause undefined behavior when the Ice run time tries to next use the factory.

The run time guarantees that `destroy` will be the last call made on the factory, that is, `create` will not be called concurrently with `destroy`, and `create` will not be called once `destroy` has been called.

Note that you cannot register a factory for the same type ID twice: if you call `addObjectFactory` with a type ID for which a factory is registered, the Ice run time throws an `AlreadyRegisteredException`.

Finally, keep in mind that if a class has only data members, but no operations, you need not create and register an object factory to transmit instances of such a class. Only if a class has operations do you have to define and register an object factory.

See Also

- Classes
- Type IDs
- Optional Data Members

Copyright 2016, ZeroC, Inc.

slice2php Command-Line Options

On this page:

- slice2php Command-Line Options
- Compiler Output in PHP
- Include Files in PHP

slice2php Command-Line Options

The Slice-to-PHP compiler, `slice2php`, offers the following command-line options in addition to the standard options:

- `--all`
  Generate code for all Slice definitions, including those included by the main Slice file.
- `-n, --namespace`
  Generate code using PHP namespaces. Note that namespaces are only supported in PHP 5.3 or later. Also note that the Ice extension for PHP must be built with namespace support enabled.
- `--checksum`
  Generate checksums for Slice definitions.

Compiler Output in PHP

For each Slice file `X.ice`, `slice2php` generates PHP code into a file named `X.php` in the output directory. The default output directory is the current working directory, but a different directory can be specified using the `--output-dir` option.

Include Files in PHP

It is important to understand how `slice2php` handles include files. In the absence of the `--all` option, the compiler does not generate PHP code for Slice definitions in included files. Rather, the compiler translates Slice `#include` statements into PHP `require` statements in the following manner:

1. Determine the full pathname of the included file.
2. Create the shortest possible relative pathname for the included file by iterating over each of the include directories (specified using the `-I` option) and removing the leading directory from the included file if possible.
   For example, if the full pathname of an included file is `/opt/App/slice/OS/Process.ice`, and we specified the options `-I/opt/App` and `-I/opt/App/slice`, then the shortest relative pathname is `OS/Process.ice` after removing `/opt/App/slice`.
3. Replace the `.ice` extension with `.php`. Continuing our example from the previous step, the translated `require` statement becomes

   ```
   require "OS/Process.php";
   ```

   As a result, you can use `-I` options to tailor the `require` statements generated by the compiler in order to avoid absolute path names and match the organizational structure of your application's source files.

See Also

- Using the Slice Compilers

Copyright 2016, ZeroC, Inc.

Application Notes for PHP

On this page:

- PHP Request Semantics
- Using Communicators in PHP
- Managing Property Sets in PHP
  - Default Property Set in PHP
  - Profiles in PHP
  - Using Property Sets in PHP
  - Security Considerations for Property Sets in PHP
- Timeouts in PHP
- Registered Communicators in PHP
  - Limitations of Registered Communicators in PHP
  - Using Registered Communicators in PHP
  - Security Considerations for Registered Communicators in PHP
  - Lifetime of Object Factories in PHP

PHP Request Semantics

In PHP terminology, a *request* is the execution of a PHP script on behalf of a Web client. Each request essentially runs in its own instance of the PHP interpreter, isolated from any other requests that may be executing concurrently. Upon the completion of a request, the interpreter reclaims memory and other resources that were acquired during the request, including objects created by the Ice extension.

Using Communicators in PHP

A communicator represents an instance of the Ice run time. A PHP script that needs to invoke an operation on a remote Ice object must initialize a communicator, obtain and narrow a proxy, and make the invocation. For example, here is a minimal (but complete) Ice script:

Copyright 2016, ZeroC, Inc.

PHP

```
<?php
require 'Ice.php';
require 'Hello.php';

$communicator = null;

try
{
    $data = new Ice_InitializationData;
    $data->properties = Ice_createProperties();
    $data->properties->load("props.cfg");
    $communicator = Ice_initialize($data);
    $proxy = $communicator->stringToProxy("...");
    $hello = Demo_HelloPrxHelper::checkedCast($proxy);
    $hello->sayHello();
}
catch(Ice_LocalException $ex)
{
    // Deal with exception...
}

if($communicator)
{
    try
    {
        $communicator->destroy();
    }
    catch(Ice_LocalException $ex)
    {
        // Ignore.
    }
}
?>
```

By default, the Ice extension automatically destroys any communicator that was created during a request. This means a script can usually omit the call to `destroy` unless there is an application-specific reason to destroy the communicator explicitly. Consequently, we can simplify our script to the following:

Copyright 2016, ZeroC, Inc.

PHP

```
<?php
require 'Ice.php';
require 'Hello.php';

try
{
    $data = new Ice_InitializationData;
    $data->properties = Ice_createProperties();
    $data->properties->load("props.cfg");
    $communicator = Ice_initialize($data);
    $proxy = $communicator->stringToProxy("...");
    $hello = Demo_HelloPrxHelper::checkedCast($proxy);
    $hello->sayHello();
}
catch(Ice_LocalException $ex)
{
    // Deal with exception...
}
?>
```

Now we allow the Ice extension to destroy our communicator automatically. (The extension traps and ignores any exception raised by `destroy`.)

Although the automatic destruction of communicators is convenient, it is important to consider the performance characteristics of this script. Specifically, each execution of the script involves the following activities:

1. Create an Ice property set
2. Load and parse a property file
3. Initialize a communicator with the given configuration properties
4. Obtain a proxy for the remote Ice object
5. Establish a socket connection to the server
6. Send a request message and wait for the reply
7. Destroy the communicator, which closes the socket connection

Of primary concern are the activities that involve system calls, such as opening and reading files, creating and using sockets, and so on. The overhead incurred by these calls may not matter if the script is only executed infrequently, but for an application with high request rates it is necessary to minimize this overhead:

- A pre-configured property set eliminates the need to parse a property file in each request.
- Timeouts prevent a script from blocking indefinitely in case Ice encounters delays while performing socket operations.
- Registering a communicator avoids the need to create and destroy a communicator in every request.
- Be aware of the number of "round trips" (request-reply pairs) your script makes. For example, the script above uses `checkedCast` to verify that the remote Ice object supports the desired Slice interface. However, calling `checkedCast` causes the Ice run time to send a request to the server and await its reply, therefore this script is actually making two remote invocations. It is unnecessary to perform a checked cast if it is safe for the client to assume that the Ice object supports the correct interface, in which case using an `uncheckedCast` instead avoids the extra round trip.

Managing Property Sets in PHP

A PHP application can manually construct a property set for configuring its communicator. The Ice extension also provides a PHP-specific property set API that helps to minimize the overhead associated with initializing a communicator, allowing you to configure a default property set along with an unlimited number of named property sets (or *profiles*). You can populate a property set using a configuration file, command-line options, or both. Property sets are initialized using the normal Ice semantics: command-line options override any settings from a configuration file.

The Ice extension creates these property sets during web server startup, which means any subsequent changes you might make to the

Copyright 2016, ZeroC, Inc.

configuration have no effect until the web server is restarted. Also keep in mind that specifying a relative path name for a configuration file usually means the path name is evaluated relative to the web server's working directory.

Default Property Set in PHP

The INI directives `ice.config` and `ice.options` specify the configuration file and the command-line options for the default property set, respectively. These directives must appear in PHP's configuration file, which is usually named `php.ini`:

```
; Snippet from php.ini on Linux
extension=IcePHP.so
ice.config=/opt/MyApp/default.cfg
ice.options="--Ice.Override.Timeout=2000"
```

Profiles in PHP

Profiles are useful when several unrelated applications execute in the same web server, or when a script needs to choose among multiple configurations. To configure your profiles, add an `ice.profiles` directive to PHP's configuration file. The value of this directive is a file containing profile definitions:

```
; Snippet from php.ini on Linux
ice.profiles=/opt/MyApp/profiles
```

The profile definition file uses INI syntax:

```
[Production]
config=/opt/MyApp/prod.cfg
options="..."

[Debug]
config=/opt/MyApp/debug.cfg
options="--Ice.Trace.Network=3 ..."
```

The name of each profile is enclosed in square brackets. The configuration file and command-line options for each profile are defined using the `config` and `options` entries, respectively.

Using Property Sets in PHP

The `Ice_getProperties` function allows a script to obtain a copy of a property set. When called without an argument, or with an empty string, the function returns the default property set. Otherwise, the function expects the name of a configured profile and returns the property set associated with that profile. The return value is an instance of `Ice_Properties`, or `null` if no matching profile was found.

Note that the Ice extension always creates the default property set, which is empty if the `ice.config` and `ice.options` directives are not defined. Also note that changes a script might make to a property set returned by this function have no effect on other requests because the script is modifying a *copy* of the original property set.

Now we can modify our script to use `Ice_getProperties` and avoid the need to load a configuration file in each request:

Copyright 2016, ZeroC, Inc.

PHP

```
<?php
require 'Ice.php';
require 'Hello.php';

try
{
    $data = new Ice_InitializationData;
    $data->properties = Ice_getProperties();
    $communicator = Ice_initialize($data);
    $proxy = $communicator->stringToProxy("...");
    $hello = Demo_HelloPrxHelper::checkedCast($proxy);
    $hello->sayHello();
}
catch(Ice_LocalException $ex)
{
    // Deal with exception...
}
?>
```

Security Considerations for Property Sets in PHP

Ice configuration properties may contain sensitive information such as the path name of the private key for an X.509 certificate. If multiple untrusted PHP applications run in the same web server, avoid the use of the default property set and choose sufficiently unique names for your named profiles. The Ice extension does not provide a means for enumerating the names of the configured profiles, therefore a malicious script would have to guess the name of a profile in order to examine its configuration properties.

To prevent a script from using the value of `ice.profiles` to open the profile definition file directly, enable the `ice.hide_profiles` direct ive to cause the Ice extension to replace the `ice.profiles` setting after it has processed the file. The `ice.hide_profiles` directive is enabled by default.

Timeouts in PHP

All twoway remote invocations made by a PHP script have synchronous semantics: the script does not regain control until Ice receives a reply from the server. As a result, we recommend configuring a suitable timeout value for all of your proxies as a defensive measure against network delays.

Registered Communicators in PHP

You can register a communicator to prevent it from being destroyed at the completion of a script. For example, a session-based PHP application can create a communicator for each new session and register it for reuse in subsequent requests of the same session. Reusing a communicator in this way avoids the overhead associated with creating and destroying a communicator in each request. Furthermore, it allows socket connections established by the Ice run time to remain open and available for use in another request.

Limitations of Registered Communicators in PHP

A communicator object is local to the process that created it, which in the case of PHP is usually a web server process. The usefulness of a registered communicator is therefore limited to situations in which an application can ensure that subsequent page requests are handled by the same web server process as the one that originally created the registered communicator. For example, registered communicators would not be appropriate in a typical CGI configuration because the CGI process terminates at the end of each request. A simple (but often impractical) solution is to configure your web server to use a single persistent process. The topic of configuring a web server to take advantage of registered communicators is outside the scope of this manual.

Using Registered Communicators in PHP

Copyright 2016, ZeroC, Inc.

The API for registered communicators consists of three functions:

- `Ice_register($communicator, $name, $expires=0)`
  Registers a communicator with the given name. On success, the function returns true. If another communicator is already registered with the same name, the function returns false. The `expires` argument specifies a timeout value in minutes; if `expires` is greater than zero, the Ice extension automatically destroys the communicator if it has not been retrieved (via `Ice_find`) for the specified number of minutes. The default value (zero) means the communicator never expires, in which case the Ice extension only destroys the communicator when the current process terminates.
  It is legal to register a communicator with more than one name. In that case, the most recent value of expires takes precedence.
- `Ice_unregister($name)`
  Removes the registration for a communicator with the given name. Returns true if a match was found or false otherwise. Calling `Ice_unregister` does not cause the communicator to be destroyed; rather, the communicator is destroyed as soon as all pending requests that are currently using the communicator have completed. Destroying a registered communicator explicitly also removes its registration.
- `Ice_find($name)`
  Retrieves the communicator associated with the given name. Returns `null` if no match is found.

An application typically uses registered communicators as follows:

PHP

```
<?php
require 'Ice.php';

$communicator = Ice_find('MyCommunicator');
$expires = ...;
if($communicator == null)
{
    $communicator = Ice_initialize(...);
    Ice_register($communicator, 'MyCommunicator', $expires);
}

...
?>
```

Note that communicators consume resources such as threads, sockets, and memory, therefore an application should be designed to minimize the number of communicators it registers. Using a suitable expiration timeout prevents registered communicators from accumulating indefinitely.

A simple application that demonstrates the use of registered communicators can be found in the `Glacier2/hello` subdirectory of the PHP sample programs.

Security Considerations for Registered Communicators in PHP

There are risks associated with allowing untrusted applications to gain access to a registered communicator. For example, if a malicious script obtains a registered communicator that is configured with SSL credentials, the script could potentially make secure invocations as if it were the trusted script.

Registering a communicator with a sufficiently unique name reduces the chance that a malicious script could guess the communicator's name. For applications that make use of PHP's session facility, the session ID is a reasonable choice for a communicator name. The sample application in `Glacier2/hello` demonstrates this solution.

Lifetime of Object Factories in PHP

PHP reclaims all memory at the end of each request, which means any object factories that a script might have installed in a registered communicator are destroyed when the request completes even if the communicator is not destroyed. As a result, a script must install its object factories in a registered communicator for *every* request, as shown in the example below:

Copyright 2016, ZeroC, Inc.

PHP

```
<?php
require 'Ice.php';

$communicator = Ice_find('MyCommunicator');
$expires = ...;
if($communicator == null)
{
    $communicator = Ice_initialize(...);
    Ice_register($communicator, 'MyCommunicator', $expires);
}

$communicator->addObjectFactory(new MyFactory, MyClass::ice_staticId());
...
?>
```

The Ice extension invokes the `destroy` method of each factory prior to the completion of a request.

See Also

- Properties and Configuration
- Connection Timeouts

Copyright 2016, ZeroC, Inc.

Using Slice Checksums in PHP

The Slice compilers can optionally generate checksums of Slice definitions. For `slice2php`, the `--checksum` option causes the compiler to generate code that adds checksums to the global array `Ice_sliceChecksums`. The checksums are installed automatically when the PHP code is first parsed; no action is required by the application.

In order to verify a server's checksums, a client could simply compare the two array objects using a comparison operator. However, this is not feasible if it is possible that the server might return a superset of the client's checksums. A more general solution is to iterate over the local checksums as demonstrated below:

PHP

```
global $Ice_sliceChecksums;
$serverChecksums = ...
foreach($Ice_sliceChecksums as $key => $value)
{
    if(!isset($serverChecksums[$key]))
        // No match found for type id!
    elseif($Ice_sliceChecksums[$key] != $serverChecksums[$key])
        // Checksum mismatch!
}
```

In this example, the client first verifies that the server's dictionary contains an entry for each Slice type ID, and then it proceeds to compare the checksums.

See Also

- Slice Checksums

Copyright 2016, ZeroC, Inc.

Example of a File System Client in PHP

This page presents a very simple client to access a server that implements the file system we developed in Slice for a Simple File System. The PHP code shown here hardly differs from the code you would write for an ordinary PHP program. This is one of the biggest advantages of using Ice: accessing a remote object is as easy as accessing an ordinary, local PHP object. This allows you to put your effort where you should, namely, into developing your application logic instead of having to struggle with arcane networking APIs.

We now have seen enough of the client-side PHP mapping to develop a complete client to access our remote file system. For reference, here is the Slice definition once more:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };

    exception GenericError {
        string reason;
    };

    sequence<string> Lines;

    interface File extends Node {
        idempotent Lines read();
        idempotent void write(Lines text) throws GenericError;
    };

    sequence<Node*> NodeSeq;

    interface Directory extends Node {
        idempotent NodeSeq list();
    };
};
```

To exercise the file system, the client does a recursive listing of the file system, starting at the root directory. For each node in the file system, the client shows the name of the node and whether that node is a file or directory. If the node is a file, the client retrieves the contents of the file and prints them.

The body of the client code looks as follows:

PHP

```
<?php
require 'Ice.php';
require 'Filesystem.php';

// Recursively print the contents of directory "dir"
// in tree fashion. For files, show the contents of
// each file. The "depth" parameter is the current
// nesting level (for indentation).

function listRecursive($dir, $depth = 0)
{
```


Copyright 2016, ZeroC, Inc.

```
    $indent = str_repeat("\t", ++$depth);

    $contents = $dir->_list(); // list is a reserved word in PHP

    foreach ($contents as $i) {
        $dir = Filesystem_DirectoryPrxHelper::checkedCast($i);
        $file = Filesystem_FilePrxHelper::uncheckedCast($i);
        echo $indent . $i->name() . ($dir ? " (directory):" : " (file):") . "\n";
        if ($dir) {
            listRecursive($dir, $depth);
        } else {
            $text = $file->read();
            foreach ($text as $j)
                echo $indent . "\t" . $j . "\n";
        }
    }
}

$ic = null;
try
{
    // Create a communicator
    //
    $ic = Ice_initialize();

    // Create a proxy for the root directory
    //
    $obj = $ic->stringToProxy("RootDir:default -p 10000");

    // Down-cast the proxy to a Directory proxy
    //
    $rootDir = Filesystem_DirectoryPrxHelper::checkedCast($obj);

    // Recursively list the contents of the root directory
    //
    echo "Contents of root directory:\n";
    listRecursive($rootDir);
}
catch(Ice_LocalException $ex)
{
    print_r($ex);
}

if($ic)
{
    // Clean up
    //
    try
    {
        $ic->destroy();
    }
    catch(Exception $ex)
    {
        print_r($ex);
    }
```


Copyright 2016, ZeroC, Inc.

```
}
?>
```

The program first defines the `listRecursive` function, which is a helper function to print the contents of the file system, and the main program follows. Let us look at the main program first:

1. The client first creates a proxy to the root directory of the file system. For this example, we assume that the server runs on the local host and listens using the default protocol (TCP/IP) at port 10000. The object identity of the root directory is known to be `RootDir`.
2. The client down-casts the proxy to the `Directory` interface and passes that proxy to `listRecursive`, which prints the contents of the file system.

Most of the work happens in `listRecursive`. The function is passed a proxy to a directory to list, and an indent level. (The indent level increments with each recursive call and allows the code to print the name of each node at an indent level that corresponds to the depth of the tree at that node.) `listRecursive` calls the `list` operation on the directory and iterates over the returned sequence of nodes:

1. The code uses `checkedCast` to narrow the `Node` proxy to a `Directory` proxy, and uses `uncheckedCast` to narrow the `Node` proxy to a `File` proxy. Exactly one of those casts will succeed, so there is no need to call `checkedCast` twice: if the `Node` *is-a* `Directory`, the code uses the proxy returned by `checkedCast`; if `checkedCast` fails, we *know* that the Node *is-a* File and, therefore, `uncheckedCast` is sufficient to get a `File` proxy.
   In general, if you know that a down-cast to a specific type will succeed, it is preferable to use `uncheckedCast` instead of `checkedCast` because `uncheckedCast` does not incur any network traffic.
2. The code prints the name of the file or directory and then, depending on which cast succeeded, prints `"(directory)"` or `"(file)"` following the name.
3. The code checks the type of the node:
   - If it is a directory, the code recurses, incrementing the indent level.
   - If it is a file, the code calls the `read` operation on the file to retrieve the file contents and then iterates over the returned sequence of lines, printing each line.

Assume that we have a small file system consisting of a two files and a a directory as follows:

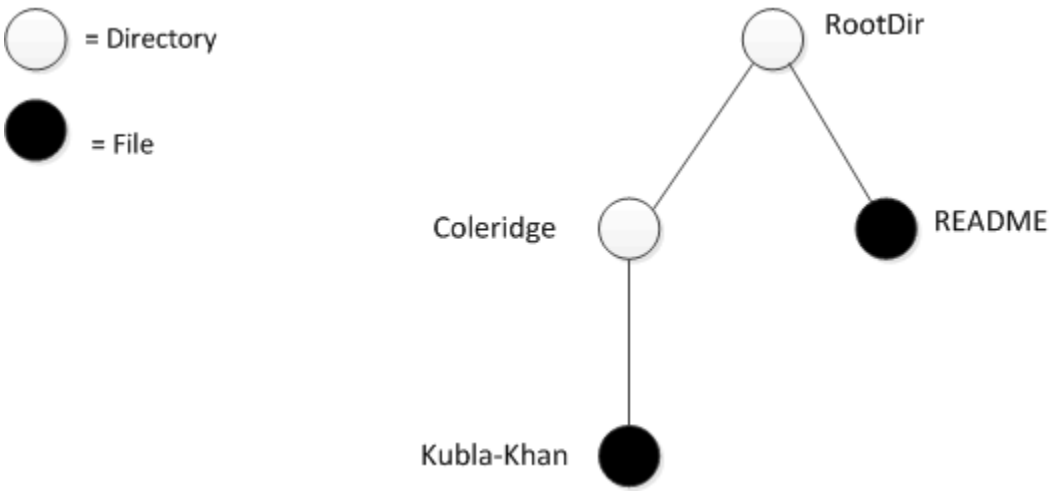

A small file system.

The output produced by the client for this file system is:

Copyright 2016, ZeroC, Inc.

```
Contents of root directory:
        README (file):
                This file system contains a collection of poetry.
        Coleridge (directory):
                Kubla_Khan (file):
                        In Xanadu did Kubla Khan
                        A stately pleasure-dome decree:
                        Where Alph, the sacred river, ran
                        Through caverns measureless to man
                        Down to a sunless sea.
```

Note that, so far, our client is not very sophisticated:

- The protocol and address information are hard-wired into the code.
- The client makes more remote procedure calls than strictly necessary; with minor redesign of the Slice definitions, many of these calls can be avoided.

We will see how to address these shortcomings in our discussions of IceGrid and object life cycle.

See Also

- Hello World Application
- Slice for a Simple File System
- Object Life Cycle
- IceGrid

Copyright 2016, ZeroC, Inc.

Python Mapping

Topics

- Client-Side Slice-to-Python Mapping
- Server-Side Slice-to-Python Mapping

Copyright 2016, ZeroC, Inc.

Client-Side Slice-to-Python Mapping

The client-side Slice-to-Python mapping defines how Slice data types are translated to Python types, and how clients invoke operations, pass parameters, and handle errors. Much of the Python mapping is intuitive. For example, Slice sequences map to Python lists, so there is essentially nothing new you have to learn in order to use Slice sequences in Python.

The Python API to the Ice run time is fully thread-safe. Obviously, you must still synchronize access to data from different threads. For example, if you have two threads sharing a sequence, you cannot safely have one thread insert into the sequence while another thread is iterating over the sequence. However, you only need to concern yourself with concurrent access to your own data — the Ice run time itself is fully thread safe, and none of the Ice API calls require you to acquire or release a lock before you safely can make the call.

Much of what appears in this chapter is reference material. We suggest that you skim the material on the initial reading and refer back to specific sections as needed. However, we recommend that you read at least the mappings for exceptions, interfaces, and operations in detail because these sections cover how to call operations from a client, pass parameters, and handle exceptions.

In order to use the Python mapping, you should need no more than the Slice definition of your application and knowledge of the Python mapping rules. In particular, looking through the generated code in order to discern how to use the Python mapping is likely to be inefficient, due to the amount of detail. Of course, occasionally, you may want to refer to the generated code to confirm a detail of the mapping, but we recommend that you otherwise use the material presented here to see how to write your client-side code.

The `Ice` Module

All of the APIs for the Ice run time are nested in the `Ice` module, to avoid clashes with definitions for other libraries or applications. Some of the contents of the `Ice` module are generated from Slice definitions; other parts of the `Ice` module provide special-purpose definitions that do not have a corresponding Slice definition. We will incrementally cover the contents of the `Ice` m odule throughout the remainder of the manual.

A Python application can load the Ice run time using the `import` statement:

```
import Ice
```

If the statement executes without error, the Ice run time is loaded and available for use. You can determine the version of the Ice run time you have just loaded by calling the `stringVersion` function:

```
icever = Ice.stringVersion()
```

Topics

- Python Mapping for Identifiers
- Python Mapping for Modules
- Python Mapping for Built-In Types
- Python Mapping for Enumerations
- Python Mapping for Structures
- Python Mapping for Sequences
- Python Mapping for Dictionaries
- Python Mapping for Constants
- Python Mapping for Exceptions
- Python Mapping for Interfaces
- Python Mapping for Operations
- Python Mapping for Classes
- Asynchronous Method Invocation (AMI) in Python
- Code Generation in Python
- Using Slice Checksums in Python
- Example of a File System Client in Python

Copyright 2016, ZeroC, Inc.

Python Mapping for Identifiers

A Slice identifier maps to an identical Python identifier. For example, the Slice identifier `Clock` becomes the Python identifier `Clock`. There is one exception to this rule: if a Slice identifier is the same as a Python keyword or is an identifier reserved by the Ice run time (such as `checkedCast`), the corresponding Python identifier is prefixed with an underscore. For example, the Slice identifier `while` is mapped as `_while`.

You should try to avoid such identifiers as much as possible.

The mapping does not modify a Slice identifier that matches the name of a Python built-in function because it can always be accessed by its fully-qualified name. For example, the built-in function `hash` can also be accessed as `__builtin__.hash`.
See Also

- Lexical Rules
- Python Mapping for Modules
- Python Mapping for Built-In Types
- Python Mapping for Enumerations
- Python Mapping for Structures
- Python Mapping for Sequences
- Python Mapping for Dictionaries
- Python Mapping for Constants
- Python Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

Python Mapping for Modules

A Slice module maps to a Python module with the same name. The mapping preserves the nesting of the Slice definitions. Note that you can optionally use packages to gain further control over the generated code.
See Also

- Modules
- Python Mapping for Identifiers
- Python Mapping for Built-In Types
- Python Mapping for Enumerations
- Python Mapping for Structures
- Python Mapping for Sequences
- Python Mapping for Dictionaries
- Python Mapping for Constants
- Python Mapping for Exceptions
- Code Generation in Python

Copyright 2016, ZeroC, Inc.

Python Mapping for Built-In Types

On this page:

- Mapping of Slice Built-In Types to Python Types
- String Mapping in Python

Mapping of Slice Built-In Types to Python Types

The Slice built-in types are mapped to Python types as shown in this table:

Slice	Python
`bool`	`bool`
`short`	`int`
`int`	`int`
`long`	`long`
`float`	`double`
`double`	`double`
`string`	`string`

Although Python supports arbitrary precision in its integer types, the Ice run time validates integer values to ensure they have valid ranges for their declared Slice types.

String Mapping in Python

String values returned as the result of a Slice operation (including return values, out parameters, and data members) are always represented as instances of Python's 8-bit `string` type. These string values contain UTF-8 encoded strings unless the program has installed a string converter, in which case string values use the converter's native encoding instead.

Legal string input values for a remote Slice operation are shown below:

- `None`
  Ice marshals an empty string whenever `None` is encountered.
- 8-bit string objects
  Ice assumes that all 8-bit string objects contain valid UTF-8 encoded strings unless the program has installed a string converter, in which case Ice assumes that 8-bit string objects use the native encoding expected by the converter.
- Unicode objects
  Ice converts a Unicode object into UTF-8 and marshals it directly. If a string converter is installed, it is not invoked for Unicode objects.

See Also

- Basic Types
- Python Mapping for Identifiers
- Python Mapping for Modules
- Python Mapping for Enumerations
- Python Mapping for Structures
- Python Mapping for Sequences
- Python Mapping for Dictionaries
- Python Mapping for Constants
- Python Mapping for Exceptions
- C++ Strings and Character Encoding

Copyright 2016, ZeroC, Inc.

Python Mapping for Enumerations

Python does not have an enumerated type, so a Slice enumeration is emulated using a Python class: the name of the Slice enumeration becomes the name of the Python class; for each enumerator, the class contains an attribute with the same name as the enumerator. For example:

Slice

```
enum Fruit { Apple, Pear, Orange };
```

The generated Python class looks as follows:

Python

```
class Fruit(Ice.EnumBase):
    def valueOf(self, n):
        # ...
    valueOf = classmethod(valueOf)

    # ...

Fruit.Apple = ...
Fruit.Pear = ...
Fruit.Orange = ...
```

Each instance of the class has a `value` attribute providing the Slice value of the enumerator, and a `name` attribute that returns its name. The `valueOf` class method translates a Slice value into its corresponding enumerator, or returns `None` if no match is found.

Given the above definitions, we can use enumerated values as follows:

Copyright 2016, ZeroC, Inc.

Python

```
f1 = Fruit.Apple
f2 = Fruit.Orange

if f1 == Fruit.Apple:               # Compare with constant
    # ...

if f1 == f2:                        # Compare two enums
    # ...

if f2.value == Fruit.Apple.value:   # Use Slice values
    # ...
elif f2.value == Fruit.Pear.value:
    # ...
elif f2.value == Fruit.Orange.value:
    # ...

Fruit.valueOf(1)  # Pear
```

As you can see, the generated class enables natural use of enumerated values. The `Fruit` class attributes are preinitialized enumerators that you can use for initialization and comparison.

Note that the generated class also defines a number of Python special methods, such as `__str__` and rich comparison operators, which we have not shown. The rich comparison operators compare the Slice value of the enumerator, which is not necessarily the same as its ordinal value.

Suppose we modify the Slice definition to include a custom enumerator value:

Slice

```
enum Fruit { Apple, Pear = 3, Orange };
```

We can use `valueOf` to examine the Slice values of the enumerators:

Python

```
Fruit.valueOf(0)  # Apple
Fruit.valueOf(1)  # None
Fruit.valueOf(3)  # Pear
Fruit.valueOf(4)  # Orange
```

See Also

- Enumerations
- Python Mapping for Identifiers
- Python Mapping for Modules
- Python Mapping for Built-In Types
- Python Mapping for Structures

Copyright 2016, ZeroC, Inc.

- Python Mapping for Sequences
- Python Mapping for Dictionaries
- Python Mapping for Constants
- Python Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

Python Mapping for Structures

A Slice structure maps to a Python class with the same name. For each Slice data member, the Python class contains a corresponding attribute. For example, here is our `Employee` structure once more:

Slice

```
struct Employee {
    long number;
    string firstName;
    string lastName;
};
```

The Python mapping generates the following definition for this structure:

Python

```
class Employee(object):
    def __init__(self, number=0, firstName='', lastName=''):
        self.number = number
        self.firstName = firstName
        self.lastName = lastName

    def __eq__(self, other):
        # ...

    def __ne__(self, other):
        # ...

    def __str__(self):
        # ...

    def __hash__(self):
        # ...

    # ...
```

The constructor initializes each of the attributes to a default value appropriate for its type:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	None

Copyright 2016, ZeroC, Inc.

`dictionary`	None
`class/interface`	None

You can also declare different default values for members of primitive and enumerated types.

The `__eq__` method returns true if all members of two structures are (recursively) equal, and `__ne__` returns true if any member differs.

The `__str__` method returns a string representation of the structure.

For structures that are also legal dictionary key types, the mapping also generates relational operators (`__lt__`, `__le__`, `__gt__`, `__ge__`) and a `__hash__` method. The `__hash__` method returns a hash value for the structure based on the value of all its data members.

See Also

- Structures
- Dictionaries
- Python Mapping for Identifiers
- Python Mapping for Modules
- Python Mapping for Built-In Types
- Python Mapping for Enumerations
- Python Mapping for Sequences
- Python Mapping for Dictionaries
- Python Mapping for Constants
- Python Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

Python Mapping for Sequences

On this page:

- Default Sequence Mapping in Python
- Allowable Sequence Values in Python
- Customizing the Sequence Mapping in Python

Default Sequence Mapping in Python

A Slice sequence maps by default to a Python list; the only exception is a sequence of bytes, which maps by default to a string in order to lower memory utilization and improve throughput. This use of native types means that the Python mapping does not generate a separate named type for a Slice sequence. It also means that you can take advantage of all the inherent functionality offered by Python's native types. For example, here is the definition of our `FruitPlatter` sequence once more:

Python

```
sequence<Fruit> FruitPlatter;
```

We can use the `FruitPlatter` sequence as shown below:

Python

```
platter = [ Fruit.Apple, Fruit.Pear ]
assert(len(platter) == 2)
platter.append(Fruit.Orange)
```

The Ice run time validates the elements of a tuple or list to ensure that they are compatible with the declared type; a `ValueError` exception is raised if an incompatible type is encountered.

Allowable Sequence Values in Python

Although each sequence type has a default mapping, the Ice run time allows a sender to use other types as well. Specifically, a tuple is also accepted for a sequence type that maps to a list, and in the case of a byte sequence, the sender is allowed to supply a tuple or list of integers as an alternative to a string.

Using a string for a byte sequence bypasses the validation step and avoids an extra copy, resulting in much greater throughput than a tuple or list. For larger byte sequences, the use of a string is strongly recommended.

Furthermore, the Ice run time accepts objects that implement Python's buffer protocol as legal values for sequences of all primitive types except strings. For example, you can use the `array` module to create a buffer that is transferred much more efficiently than a tuple or list. Consider the two sequence values in the sample code below:

Python

```
import array
...
seq1 = array.array("i", [1, 2, 3, 4, 5])
seq2 = [1, 2, 3, 4, 5]
```

The values have the same on-the-wire representation, but they differ greatly in marshaling overhead because the buffer can be traversed

Copyright 2016, ZeroC, Inc.

more quickly and requires no validation.

Note that the Ice run time has no way of knowing what type of elements a buffer contains, therefore it is the application's responsibility to ensure that a buffer is compatible with the declared sequence type.

Customizing the Sequence Mapping in Python

The previous section described the allowable types that an application may use when sending a sequence. That kind of flexibility is not possible when receiving a sequence, because in this case it is the Ice run time's responsibility to create the container that holds the sequence.

As stated earlier, the default mapping for most sequence types is a list, and for byte sequences the default mapping is a string. Unless otherwise indicated, an application always receives sequences as the container type specified by the default mapping. If it would be more convenient to receive a sequence as a different type, you can customize the mapping by annotating your Slice definitions with metadata. The following table describes the metadata directives supported by the Python mapping:

Directive	Description
`python:seq:default`	Use the default mapping.
`python:seq:list`	Map to a Python list.
`python:seq:tuple`	Map to a Python tuple.

A metadata directive may be specified when defining a sequence, or when a sequence is used as a parameter, return value or data member. If specified at the point of definition, the directive affects all occurrences of that sequence type unless overridden by another directive at a point of use. The following Slice definitions illustrate these points:

Slice

```
sequence<int> IntList; // Uses list by default
["python:seq:tuple"] sequence<int> IntTuple; // Defaults to tuple

sequence<byte> ByteString; // Uses string by default
["python:seq:list"] sequence<byte> ByteList; // Defaults to list

struct S {
    IntList i1; // list
    IntTuple i2; // tuple
    ["python:seq:tuple"] IntList i3; // tuple
    ["python:seq:list"] IntTuple i4; // list
    ["python:seq:default"] IntTuple i5; // list

    ByteString b1; // string
    ByteList b2; // list
    ["python:seq:list"] ByteString b3; // list
    ["python:seq:tuple"] ByteString b4; // tuple
    ["python:seq:default"] ByteList b5; // string
};

interface I {
    IntList op1(ByteString s1, out ByteList s2);

    ["python:seq:tuple"]
    IntList op2(["python:seq:list"] ByteString s1,
                ["python:seq:tuple"] out ByteList s2);
};
```


Copyright 2016, ZeroC, Inc.

The operation `op2` and the data members of structure `S` demonstrate how to override the mapping for a sequence at the point of use.

It is important to remember that these metadata directives only affect the receiver of the sequence. For example, the data members of structure `S` are populated with the specified sequence types only when the Ice run time unmarshals an instance of `S`. In the case of an operation, custom metadata affects the client when specified for the operation's return type and output parameters, whereas metadata affects the server for input parameters.

See Also

- Sequences
- Python Mapping for Identifiers
- Python Mapping for Modules
- Python Mapping for Built-In Types
- Python Mapping for Enumerations
- Python Mapping for Structures
- Python Mapping for Dictionaries
- Python Mapping for Constants
- Python Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

Python Mapping for Dictionaries

Here is the definition of our `EmployeeMap` once more:

Slice

```
dictionary<long, Employee> EmployeeMap;
```

As for sequences, the Python mapping does not create a separate named type for this definition. Instead, *all* dictionaries are simply instances of Python's dictionary type. For example:

Python

```
em = {}

e = Employee()
e.number = 31
e.firstName = "James"
e.lastName = "Gosling"

em[e.number] = e
```

The Ice run time validates the elements of a dictionary to ensure that they are compatible with the declared type; a `ValueError` exception is raised if an incompatible type is encountered.
See Also

- Dictionaries
- Python Mapping for Identifiers
- Python Mapping for Modules
- Python Mapping for Built-In Types
- Python Mapping for Enumerations
- Python Mapping for Structures
- Python Mapping for Sequences
- Python Mapping for Constants
- Python Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

Python Mapping for Constants

Here are the constant definitions once more:

Slice

```
const bool      AppendByDefault = true;
const byte      LowerNibble = 0x0f;
const string    Advice = "Don't Panic!";
const short     TheAnswer = 42;
const double    PI = 3.1416;

enum Fruit { Apple, Pear, Orange };
const Fruit     FavoriteFruit = Pear;
```

The generated definitions for these constants are shown below:

Python

```
AppendByDefault = True
LowerNibble = 15
Advice = "Don't Panic!"
TheAnswer = 42
PI = 3.1416
FavoriteFruit = Fruit.Pear
```

As you can see, each Slice constant is mapped to a Python attribute with the same name as the constant.

Slice string literals that contain non-ASCII characters or universal character names are mapped to Python string literals with these characters replaced by their UTF-8 encoding as octal escapes. For example:

Slice

```
const string Egg = "œuf";
const string Heart = "c\u0153ur";
const string Banana = "\U0001F34C";
```

is mapped to:

Python

```
Egg = "\305\223uf"
Heart = "c\305\223ur"
Banana = "\360\237\215\214"
```

See Also

- Constants and Literals
- Python Mapping for Identifiers
- Python Mapping for Modules
- Python Mapping for Built-In Types

Copyright 2016, ZeroC, Inc.

- Python Mapping for Enumerations
- Python Mapping for Structures
- Python Mapping for Sequences
- Python Mapping for Dictionaries
- Python Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

Python Mapping for Exceptions

On this page:

- Inheritance Hierarchy for Exceptions in Python
- Python Mapping for User Exceptions
  - Optional Data Members
- Python Mapping for Run-Time Exceptions

Inheritance Hierarchy for Exceptions in Python

The mapping for exceptions is based on the inheritance hierarchy shown below:

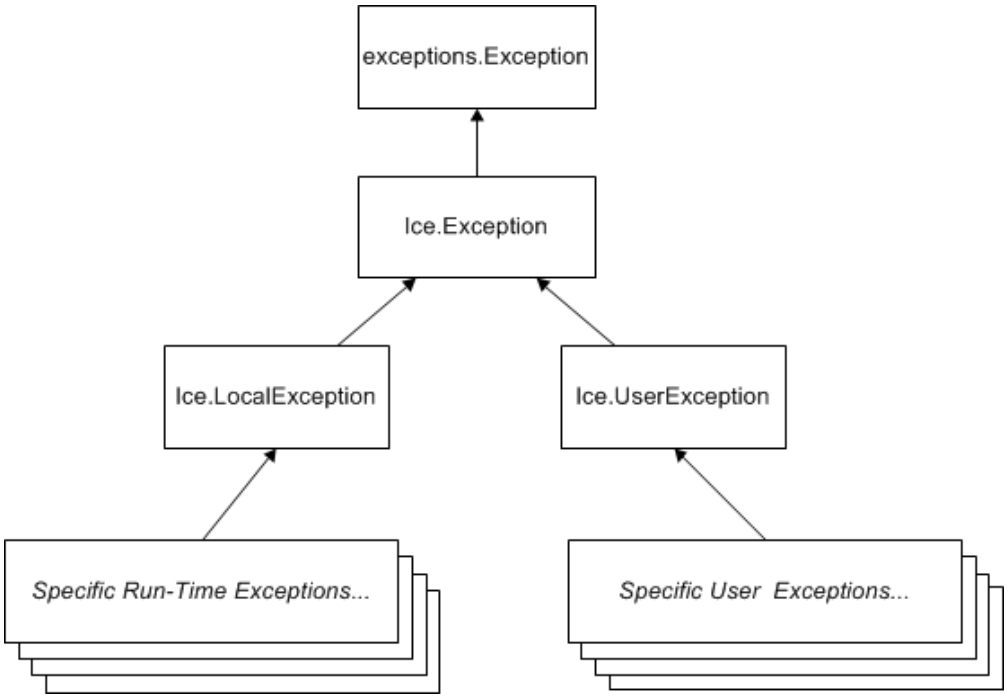

Inheritance structure for Ice exceptions.

The ancestor of all exceptions is `exceptions.Exception`, from which `Ice.Exception` is derived. `Ice.LocalException` and `Ice.UserException` are derived from `Ice.Exception` and form the base for all run-time and user exceptions.

Python Mapping for User Exceptions

Here is a fragment of the Slice definition for our world time server once more:

Slice

```
exception GenericError {
    string reason;
};
exception BadTimeVal extends GenericError {};
exception BadZoneName extends GenericError {};
```

These exception definitions map as follows:

Copyright 2016, ZeroC, Inc.

Python

```
class GenericError(Ice.UserException):
    def __init__(self, reason=''):
        self.reason = reason

    def ice_name(self):
        # ...

    def __str__(self):
        # ...

class BadTimeVal(GenericError):
    def __init__(self, reason=''):
        GenericError.__init__(self, reason)

    def ice_name(self):
        # ...

    def __str__(self):
        # ...

class BadZoneName(GenericError):
    def __init__(self, reason=''):
        GenericError.__init__(self, reason)

    def ice_name(self):
        # ...

    def __str__(self):
        # ...
```

Each Slice exception is mapped to a Python class with the same name. The inheritance structure of the Slice exceptions is preserved for the generated classes, so `BadTimeVal` and `BadZoneName` inherit from `GenericError`.

Each exception member corresponds to an attribute of the instance, which the constructor initializes to a default value appropriate for its type:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	None
`dictionary`	None
`class/interface`	None

You can also declare different default values for members of primitive and enumerated types. For derived exceptions, the constructor has

Copyright 2016, ZeroC, Inc.

one parameter for each of the base exception's data members, plus one parameter for each of the derived exception's data members, in base-to-derived order. As an example, although `BadTimeVal` and `BadZoneName` do not declare data members, their constructors still accept a value for the inherited data member `reason` in order to pass it to the constructor of the base exception `GenericError`.

Each exception also defines the `ice_name` method to return the name of the exception, and the special method `__str__` to return a stringified representation of the exception and its members.

All user exceptions are derived from the base class `Ice.UserException`. This allows you to catch all user exceptions generically by installing a handler for `Ice.UserException`. Similarly, you can catch all Ice run-time exceptions with a handler for `Ice.LocalException`, and you can catch all Ice exceptions with a handler for `Ice.Exception`.

Optional Data Members

Optional data members use the same mapping as required data members, but an optional data member can also be set to the marker value `Ice.Unset` to indicate that the member is unset. A well-behaved program must test an optional data member before using its value:

Python

```
try:
    ...
except ex:
    if ex.optionalMember:
        print("optionalMember = " + str(ex.optionalMember))
    else:
        print("optionalMember is unset")
```

The `Ice.Unset` marker value has different semantics than `None`. Since `None` is a legal value for certain Slice types, the Ice run time requires a separate marker value so that it can determine whether an optional value is set. An optional value set to `None` is considered to be set. If you need to distinguish between an unset value and a value set to `None`, you can do so as follows:

Python

```
try:
    ...
except ex:
    if ex.optionalMember is Ice.Unset:
        print("optionalMember is unset")
    elif ex.optionalMember is None:
        print("optionalMember is None")
    else:
        print("optionalMember = " + str(ex.optionalMember))
```

Python Mapping for Run-Time Exceptions

The Ice run time throws run-time exceptions for a number of pre-defined error conditions. All run-time exceptions directly or indirectly derive from `Ice.LocalException` (which, in turn, derives from `Ice.Exception`).

By catching exceptions at the appropriate point in the inheritance hierarchy, you can handle exceptions according to the category of error they indicate:

- `Ice.LocalException`
  This is the root of the inheritance tree for run-time exceptions.
- `Ice.UserException`

Copyright 2016, ZeroC, Inc.

This is the root of the inheritance tree for user exceptions.

- `Ice.TimeoutException`
  This is the base exception for both operation-invocation and connection-establishment timeouts.
- `Ice.ConnectTimeoutException`
  This exception is raised when the initial attempt to establish a connection to a server times out.

For example, a `ConnectTimeoutException` can be handled as `ConnectTimeoutException`, `TimeoutException`, `LocalException`, or `Exception`.

You will probably have little need to catch run-time exceptions as their most-derived type and instead catch them as `LocalException`; the fine-grained error handling offered by the remainder of the hierarchy is of interest mainly in the implementation of the Ice run time. Exceptions to this rule are the exceptions related to facet and object life cycles, which you may want to catch explicitly. These exceptions are `FacetNotExistException` and `ObjectNotExistException`, respectively.

See Also

- User Exceptions
- Run-Time Exceptions
- Python Mapping for Identifiers
- Python Mapping for Modules
- Python Mapping for Built-In Types
- Python Mapping for Enumerations
- Python Mapping for Structures
- Python Mapping for Sequences
- Python Mapping for Dictionaries
- Python Mapping for Constants
- Optional Data Members
- Versioning
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

Python Mapping for Interfaces

The mapping of Slice interfaces revolves around the idea that, to invoke a remote operation, you call a member function on a local class instance that is a proxy for the remote object. This makes the mapping easy and intuitive to use because making a remote procedure call is no different from making a local procedure call (apart from error semantics).

On this page:

- Proxy Classes in Python
- Ice.ObjectPrx Class in Python
- Casting Proxies in Python
- Using Proxy Methods in Python
- Object Identity and Proxy Comparison in Python

Proxy Classes in Python

On the client side, a Slice interface maps to a Python class with methods that correspond to the operations on that interface. Consider the following simple interface:

Slice

```
interface Simple {
    void op();
};
```

The Python mapping generates the following definition for use by the client:

Python

```
class SimplePrx(Ice.ObjectPrx):
    def op(self, _ctx=None):
        # ...

    def ice_staticId():
        return '...'
    ice_staticId = staticmethod(ice_staticId)

    # ...
```

In the client's address space, an instance of `SimplePrx` is the local ambassador for a remote instance of the `Simple` interface in a server and is known as a *proxy instance*. All the details about the server-side object, such as its address, what protocol to use, and its object identity are encapsulated in that instance.

Note that `SimplePrx` inherits from `Ice.ObjectPrx`. This reflects the fact that all Ice interfaces implicitly inherit from `Ice::Object`.

For each operation in the interface, the proxy class has a method of the same name. In the preceding example, we find that the operation `op` has been mapped to the method `op`. Note that `op` accepts an optional trailing parameter `_ctx` representing the operation context. This parameter is a Python dictionary for use by the Ice run time to store information about how to deliver a request. You normally do not need to use it. (We examine the context parameter in detail in Request Contexts. The parameter is also used by IceStorm.)

Proxy instances are always created on behalf of the client by the Ice run time, so client code never has any need to instantiate a proxy directly.

A value of `None` denotes the null proxy. The null proxy is a dedicated value that indicates that a proxy points "nowhere" (denotes no object).

Another method defined by every proxy class is `ice_staticId`, which returns the type ID string corresponding to the interface. As an

Copyright 2016, ZeroC, Inc.

example, for the Slice interface `Simple` in module `M`, the string returned by `ice_staticId` is `"::M::Simple"`.

`Ice.ObjectPrx` Class in Python

All Ice objects have `Object` as the ultimate ancestor type, so all proxies inherit from `Ice.ObjectPrx`. `ObjectPrx` provides a number of methods:

Python

```
class ObjectPrx(object):
    def equals(self, other):
    def ice_getIdentity(self):
    def ice_isA(self, id):
    def ice_ids(self):
    def ice_id(self):
    def ice_ping(self):
    # ...
```

The methods behave as follows:

- **`equals`**
  This method compares two proxies for equality. Note that all aspects of proxies are compared by this method, such as the communication endpoints for the proxy. This means that, in general, if two proxies compare unequal, that does *not* imply that they denote different objects. For example, if two proxies denote the same Ice object via different transport endpoints, `equals` returns `false` even though the proxies denote the same object.

- **`ice_getIdentity`**
  This method returns the identity of the object denoted by the proxy. The identity of an Ice object has the following Slice type:

  Slice

  ```
  module Ice {
      struct Identity {
          string name;
          string category;
      };
  };
  ```

  To see whether two proxies denote the same object, first obtain the identity for each object and then compare the identities:

Copyright 2016, ZeroC, Inc.

Python

```
proxy1 = ...
proxy2 = ...
id1 = proxy1.ice_getIdentity()
id2 = proxy2.ice_getIdentity()

if id1 == id2:
    # proxy1 and proxy2 denote the same object
else:
    # proxy1 and proxy2 denote different objects
```

- **ice_isA**
  The ice_isA method determines whether the object denoted by the proxy supports a specific interface. The argument to ice_isA is a type ID. For example, to see whether a proxy of type ObjectPrx denotes a Printer object, we can write:

  Python

  ```
  proxy = ...
  if proxy != None and proxy.ice_isA("::Printer"):
      # proxy denotes a Printer object
  else:
      # proxy denotes some other type of object
  ```

  Note that we are testing whether the proxy is None before attempting to invoke the ice_isA method. This avoids getting a run-time error if the proxy is None.

- **ice_ids**
  The ice_ids method returns an array of strings representing all of the type IDs that the object denoted by the proxy supports.

- **ice_id**
  The ice_id method returns the type ID of the object denoted by the proxy. Note that the type returned is the type of the actual object, which may be more derived than the static type of the proxy. For example, if we have a proxy of type BasePrx, with a static type ID of ::Base, the return value of ice_id might be ::Base, or it might something more derived, such as ::Derived.

- **ice_ping**
  The ice_ping method provides a basic reachability test for the object. If the object can physically be contacted (that is, the object exists and its server is running and reachable), the call completes normally; otherwise, it throws an exception that indicates why the object could not be reached, such as ObjectNotExistException or ConnectTimeoutException.

The ice_isA, ice_ids, ice_id, and ice_ping methods are remote operations and therefore support an optional trailing parameter representing a request context. Also note that there are other methods in ObjectPrx, not shown here. These methods provide different ways to dispatch a call and also provide access to an object's facets.

Casting Proxies in Python

The Python mapping for a proxy also generates two static methods:

Copyright 2016, ZeroC, Inc.

Python

```
class SimplePrx(Ice.ObjectPrx):
    # ...

    def checkedCast(proxy, facet=''):
        # ...
    checkedCast = staticmethod(checkedCast)

    def uncheckedCast(proxy, facet=''):
        # ...
    uncheckedCast = staticmethod(uncheckedCast)
```

The method names `checkedCast` and `uncheckedCast` are reserved for use in proxies. If a Slice interface defines an operation with either of those names, the mapping escapes the name in the generated proxy by prepending an underscore. For example, an interface that defines an operation named `checkedCast` is mapped to a proxy with a method named `_checkedCast`.

For `checkedCast`, if the passed proxy is for an object of type `Simple`, or a proxy for an object with a type derived from `Simple`, the cast returns a reference to a proxy of type `SimplePrx`; otherwise, if the passed proxy denotes an object of a different type (or if the passed proxy is `None`), the cast returns `None`.

Given a proxy of any type, you can use a `checkedCast` to determine whether the corresponding object supports a given type, for example:

Python

```
obj = ...       # Get a proxy from somewhere...

simple = SimplePrx.checkedCast(obj)
if simple != None:
    # Object supports the Simple interface...
else:
    # Object is not of type Simple...
```

Note that a `checkedCast` contacts the server. This is necessary because only the implementation of an object in the server has definite knowledge of the type of an object. As a result, a `checkedCast` may throw a `ConnectTimeoutException` or an `ObjectNotExistException`.

In contrast, an `uncheckedCast` does not contact the server and unconditionally returns a proxy of the requested type. However, if you do use an `uncheckedCast`, you must be certain that the proxy really does support the type you are casting to; otherwise, if you get it wrong, you will most likely get a run-time exception when you invoke an operation on the proxy. The most likely error for such a type mismatch is `OperationNotExistException`. However, other exceptions, such as a marshaling exception are possible as well. And, if the object happens to have an operation with the correct name, but different parameter types, no exception may be reported at all and you simply end up sending the invocation to an object of the wrong type; that object may do rather nonsensical things. To illustrate this, consider the following two interfaces:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Process {
    void launch(int stackSize, int dataSize);
};

// ...

interface Rocket {
    void launch(float xCoord, float yCoord);
};
```

Suppose you expect to receive a proxy for a `Process` object and use an `uncheckedCast` to down-cast the proxy:

Python

```
obj = ...                               # Get proxy...
process = ProcessPrx.uncheckedCast(obj) # No worries...
process.launch(40, 60)                  # Oops...
```

If the proxy you received actually denotes a `Rocket` object, the error will go undetected by the Ice run time: because `int` and `float` have the same size and because the Ice protocol does not tag data with its type on the wire, the implementation of `Rocket::launch` will simply misinterpret the passed integers as floating-point numbers.

In fairness, this example is somewhat contrived. For such a mistake to go unnoticed at run time, both objects must have an operation with the same name and, in addition, the run-time arguments passed to the operation must have a total marshaled size that matches the number of bytes that are expected by the unmarshaling code on the server side. In practice, this is extremely rare and an incorrect `uncheckedCast` typically results in a run-time exception.

Using Proxy Methods in Python

The base proxy class `ObjectPrx` supports a variety of methods for customizing a proxy. Since proxies are immutable, each of these "factory methods" returns a copy of the original proxy that contains the desired modification. For example, you can obtain a proxy configured with a ten second timeout as shown below:

Python

```
proxy = communicator.stringToProxy(...)
proxy = proxy.ice_timeout(10000)
```

A factory method returns a new proxy object if the requested modification differs from the current proxy, otherwise it returns the current proxy. With few exceptions, factory methods return a proxy of the same type as the current proxy, therefore it is generally not necessary to repeat a down-cast after using a factory method. The example below demonstrates these semantics:

Copyright 2016, ZeroC, Inc.

Python

```
base = communicator.stringToProxy(...)
hello = Demo.HelloPrx.checkedCast(base)
hello = hello.ice_timeout(10000) # Type is preserved
hello.sayHello()
```

The only exceptions are the factory methods `ice_facet` and `ice_identity`. Calls to either of these methods may produce a proxy for an object of an unrelated type, therefore they return a base proxy that you must subsequently down-cast to an appropriate type.

Object Identity and Proxy Comparison in Python

Proxy objects support comparison using the built-in relational operators as well as the `cmp` function. Note that proxy comparison uses *all* of the information in a proxy for the comparison. This means that not only the object identity must match for a comparison to succeed, but other details inside the proxy, such as the protocol and endpoint information, must be the same. In other words, comparison tests for *proxy* identity , *not* object identity. A common mistake is to write code along the following lines:

Python

```
p1 = ...        # Get a proxy...
p2 = ...        # Get another proxy...

if p1 != p2:
    # p1 and p2 denote different objects       # WRONG!
else:
    # p1 and p2 denote the same object         # Correct
```

Even though `p1` and `p2` differ, they may denote the same Ice object. This can happen because, for example, both `p1` and `p2` embed the same object identity, but each uses a different protocol to contact the target object. Similarly, the protocols may be the same, but denote different endpoints (because a single Ice object can be contacted via several different transport endpoints). In other words, if two proxies compare equal, we know that the two proxies denote the same object (because they are identical in all respects); however, if two proxies compare unequal, we know absolutely nothing: the proxies may or may not denote the same object.

To compare the object identities of two proxies, you can use helper functions in the `Ice` module:

Python

```
def proxyIdentityCompare(lhs, rhs)
def proxyIdentityAndFacetCompare(lhs, rhs)
```

`proxyIdentityCompare` allows you to correctly compare proxies for identity:

Copyright 2016, ZeroC, Inc.

Python

```
p1 = ...        # Get a proxy...
p2 = ...        # Get another proxy...

if Ice.proxyIdentityCompare(p1, p2) != 0:
    # p1 and p2 denote different objects       # Correct
else:
    # p1 and p2 denote the same object         # Correct
```

The function returns 0 if the identities are equal, `1` if `p1` is less than `p2`, and `1` if `p1` is greater than `p2`. (The comparison uses `name` as the major sort key and `category` as the minor sort key.)

The `proxyIdentityAndFacetCompare` function behaves similarly, but compares both the identity and the facet name.

See Also

- Interfaces, Operations, and Exceptions
- Proxies for Ice Objects
- Python Mapping for Operations
- Operations on Object
- Proxy Methods
- Versioning
- IceStorm

Copyright 2016, ZeroC, Inc.

Python Mapping for Operations

On this page:

- Basic Python Mapping for Operations
- Normal and idempotent Operations in Python
- Passing Parameters in Python
  - In-Parameters in Python
  - Out-Parameters in Python
  - Parameter Type Mismatches in Python
  - Null Parameters in Python
  - Optional Parameters in Python
- Exception Handling in Python

Basic Python Mapping for Operations

As we saw in the Python mapping for interfaces, for each operation on an interface, the proxy class contains a corresponding method with the same name. To invoke an operation, you call it via the proxy. For example, here is part of the definitions for our file system:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };
    // ...
};
```

The `name` operation returns a value of type `string`. Given a proxy to an object of type `Node`, the client can invoke the operation as follows:

Python

```
node = ...          # Initialize proxy
name = node.name()  # Get name via RPC
```

Normal and `idempotent` Operations in Python

You can add an `idempotent` qualifier to a Slice operation. As far as the signature for the corresponding proxy method is concerned, `idempotent` has no effect. For example, consider the following interface:

Slice

```
interface Example {
                string op1();
    idempotent  string op2();
};
```

The proxy class for this is:

Copyright 2016, ZeroC, Inc.

Python

```
class ExamplePrx(Ice.ObjectPrx):
    def op1(self, _ctx=None):
        # ...

    def op2(self, _ctx=None):
        # ...
```

Because `idempotent` affects an aspect of call dispatch, not interface, it makes sense for the two methods to look the same.

Passing Parameters in Python

In-Parameters in Python

All parameters are passed by reference in the Python mapping; it is guaranteed that the value of a parameter will not be changed by the invocation.

Here is an interface with operations that pass parameters of various types from client to server:

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ClientToServer {
    void op1(int i, float f, bool b, string s);
    void op2(NumberAndString ns, StringSeq ss, StringTable st);
    void op3(ClientToServer* proxy);
};
```

The Slice compiler generates the following proxy for this definition:

Copyright 2016, ZeroC, Inc.

Python

```
class ClientToServerPrx(Ice.ObjectPrx):
    def op1(self, i, f, b, s, _ctx=None):
        # ...

    def op2(self, ns, ss, st, _ctx=None):
        # ...

    def op3(self, proxy, _ctx=None):
        # ...
```

Given a proxy to a `ClientToServer` interface, the client code can pass parameters as in the following example:

Python

```
p = ...                                 # Get proxy...

p.op1(42, 3.14f, True, "Hello world!")  # Pass simple literals

i = 42
f = 3.14f
b = True
s = "Hello world!"
p.op1(i, f, b, s)                       # Pass simple variables

ns = NumberAndString()
ns.x = 42
ns.str = "The Answer"
ss = [ "Hello world!" ]
st = {}
st[0] = ns
p.op2(ns, ss, st)                       # Pass complex variables

p.op3(p)                                # Pass proxy
```

Out-Parameters in Python

As in Java, Python functions do not support reference arguments. That is, it is not possible to pass an uninitialized variable to a Python function in order to have its value initialized by the function. The Java mapping overcomes this limitation with the use of *holder classes* that represent each `out` parameter. The Python mapping takes a different approach, one that is more natural for Python users.

The semantics of `out` parameters in the Python mapping depend on whether the operation returns one value or multiple values. An operation returns multiple values when it has declared multiple `out` parameters, or when it has declared a non-`void` return type and at least one `out` parameter.

If an operation returns multiple values, the client receives them in the form of a *result tuple*. A non-`void` return value, if any, is always the first element in the result tuple, followed by the `out` parameters in the order of declaration.

If an operation returns only one value, the client receives the value itself.

Copyright 2016, ZeroC, Inc.

Here again are the same Slice definitions we saw earlier, but this time with all parameters being passed in the `out` direction:

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ServerToClient {
    int op1(out float f, out bool b, out string s);
    void op2(out NumberAndString ns,
             out StringSeq ss,
             out StringTable st);
    void op3(out ServerToClient* proxy);
};
```

The Python mapping generates the following code for this definition:

Python

```
class ServerToClientPrx(Ice.ObjectPrx):
    def op1(self, _ctx=None):
        # ...

    def op2(self, _ctx=None):
        # ...

    def op3(self, _ctx=None):
        # ...
```

Given a proxy to a `ServerToClient` interface, the client code can receive the results as in the following example:

Python

```
p = ...              # Get proxy...
i, f, b, s = p.op1()
ns, ss, st = p.op2()
stcp = p.op3()
```

The operations have no `in` parameters, therefore no arguments are passed to the proxy methods. Since `op1` and `op2` return multiple values, their result tuples are unpacked into separate values, whereas the return value of `op3` requires no unpacking.

Parameter Type Mismatches in Python

Copyright 2016, ZeroC, Inc.

Although the Python compiler cannot check the types of arguments passed to a function, the Ice run time does perform validation on the arguments to a proxy invocation and reports any type mismatches as a `ValueError` exception.

Null Parameters in Python

Some Slice types naturally have "empty" or "not there" semantics. Specifically, sequences, dictionaries, and strings all can be `None`, but the corresponding Slice types do not have the concept of a null value. To make life with these types easier, whenever you pass `None` as a parameter or return value of type sequence, dictionary, or string, the Ice run time automatically sends an empty sequence, dictionary, or string to the receiver.

This behavior is useful as a convenience feature: especially for deeply-nested data types, members that are sequences, dictionaries, or strings automatically arrive as an empty value at the receiving end. This saves you having to explicitly initialize, for example, every string element in a large sequence before sending the sequence in order to avoid a run-time error. Note that using null parameters in this way does *not* create null semantics for Slice sequences, dictionaries, or strings. As far as the object model is concerned, these do not exist (only *empty* sequences, dictionaries, and strings do). For example, it makes no difference to the receiver whether you send a string as `None` or as an empty string: either way, the receiver sees an empty string.

Optional Parameters in Python

Optional parameters use the same mapping as required parameters. The only difference is that `Ice.Unset` can be passed as the value of an optional parameter or return value. Consider the following operation:

Slice

```
optional(1) int execute(optional(2) string params, out optional(3) float value);
```

A client can invoke this operation as shown below:

Python

```
i, v = proxy.execute("--file log.txt")
i, v = proxy.execute(Ice.Unset)

if v:
    print("value = " + str(v)) # v is set to a value
```

A well-behaved program must always test an optional parameter prior to using its value. Keep in mind that the `Ice.Unset` marker value has different semantics than `None`. Since `None` is a legal value for certain Slice types, the Ice run time requires a separate marker value so that it can determine whether an optional parameter is set. An optional parameter set to `None` is considered to be set. If you need to distinguish between an unset parameter and a parameter set to `None`, you can do so as follows:

Python

```
if optionalParam is Ice.Unset:
    print("optionalParam is unset")
elif optionalParam is None:
    print("optionalParam is None")
else:
    print("optionalParam = " + str(optionalParam))
```


Copyright 2016, ZeroC, Inc.

Exception Handling in Python

Any operation invocation may throw a run-time exception and, if the operation has an exception specification, may also throw user exceptions. Suppose we have the following simple interface:

Slice

```
exception Tantrum {
    string reason;
};

interface Child {
    void askToCleanUp() throws Tantrum;
};
```

Slice exceptions are thrown as Python exceptions, so you can simply enclose one or more operation invocations in a `try-except` block:

Python

```
child = ...       # Get child proxy...

try:
    child.askToCleanUp()
except Tantrum, t:
    print "The child says:", t.reason
```

Typically, you will catch only a few exceptions of specific interest around an operation invocation; other exceptions, such as unexpected run-time errors, will usually be handled by exception handlers higher in the hierarchy. For example:

Python

```
import traceback, Ice

def run():
    child = ...   # Get child proxy...
    try:
        child.askToCleanUp()
    except Tantrum, t:
        print "The child says:", t.reason
        child.scold() # Recover from error...
    child.praise() # Give positive feedback...

try:
    # ...
    run()
    # ...
except Ice.Exception:
    traceback.print_exc()
```


Copyright 2016, ZeroC, Inc.

This code handles a specific exception of local interest at the point of call and deals with other exceptions generically. (This is also the strategy we used for our first simple application in Hello World Application.)

See Also

- Operations
- Hello World Application
- Slice for a Simple File System
- Python Mapping for Operations
- Python Mapping for Interfaces
- Python Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

Python Mapping for Classes

On this page:

- Basic Python Mapping for Classes
- Inheritance from Ice.Object in Python
- Class Data Members in Python
- Class Constructors in Python
- Class Operations in Python
- Class Factories in Python

Basic Python Mapping for Classes

A Slice class maps to a Python class with the same name. The generated class contains an attribute for each Slice data member (just as for structures and exceptions). Consider the following class definition:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Python mapping generates the following code for this definition:

Python

```
class TimeOfDay(Ice.Object):
    def __init__(self, hour=0, minute=0, second=0):
        # ...
        self.hour = hour
        self.minute = minute
        self.second = second

    def ice_staticId():
        return '::M::TimeOfDay'
    ice_staticId = staticmethod(ice_staticId)

    # ...

    #
    # Operation signatures.
    #
    # def format(self, current=None):
```

There are a number of things to note about the generated code:

1. The generated class `TimeOfDay` inherits from `Ice.Object`. This means that all classes implicitly inherit from `Ice.Object`, which is the ultimate ancestor of all classes. Note that `Ice.Object` is *not* the same as `Ice.ObjectPrx`. In other words, you *cannot* pass a class where a proxy is expected and vice versa.
2. The constructor defines an attribute for each Slice data member.
3. The class defines the static method `ice_staticId`.

Copyright 2016, ZeroC, Inc.

4. A comment summarizes the method signatures for each Slice operation.

There is quite a bit to discuss here, so we will look at each item in turn.

Inheritance from `Ice.Object` in Python

Like interfaces, classes implicitly inherit from a common base class, `Ice.Object`. However, as shown in the illustration below, classes inherit from `Ice.Object` instead of `Ice.ObjectPrx` (which is at the base of the inheritance hierarchy for proxies). As a result, you cannot pass a class where a proxy is expected (and vice versa) because the base types for classes and proxies are not compatible.

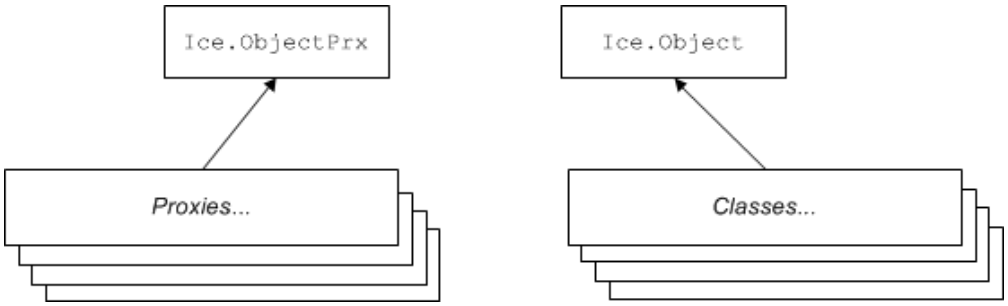

Inheritance from `Ice.ObjectPrx` and `Ice.Object`.

`Ice.Object` contains a number of methods:

Python

```
class Object(object):
    def ice_isA(self, id, current=None):
        # ...

    def ice_ping(self, current=None):
        # ...

    def ice_ids(self, current=None):
        # ...

    def ice_id(self, current=None):
        # ...

    def ice_staticId():
        # ...
    ice_staticId = staticmethod(ice_staticId)

    def ice_preMarshal(self):
        # ...

    def ice_postUnmarshal(self):
        # ...
```

The member functions of `Ice.Object` behave as follows:

- `ice_isA`
  This method returns `true` if the object supports the given type ID, and `false` otherwise.
- `ice_ping`

Copyright 2016, ZeroC, Inc.

As for interfaces, `ice_ping` provides a basic reachability test for the object.

- `ice_ids`
  This method returns a string sequence representing all of the type IDs supported by this object, including `::Ice::Object`.
- `ice_id`
  This method returns the actual run-time type ID of the object. If you call `ice_id` through a reference to a base instance, the returned type ID is the actual (possibly more derived) type ID of the instance.
- `ice_staticId`
  This method is generated in each class and returns the static type ID of the class.
- `ice_preMarshal`
  The Ice run time invokes this method prior to marshaling the object's state, providing the opportunity for a subclass to validate its declared data members.
- `ice_postUnmarshal`
  The Ice run time invokes this method after unmarshaling an object's state. A subclass typically overrides this function when it needs to perform additional initialization using the values of its declared data members.

Note that neither `Ice.Object` nor the generated class override `__hash__` and `__eq__`, so the default implementations apply.

Class Data Members in Python

By default, data members of classes are mapped exactly as for structures and exceptions: for each data member in the Slice definition, the generated class contains a corresponding attribute.

Optional data members use the same mapping as required data members, but an optional data member can also be set to the marker value `Ice.Unset` to indicate that the member is unset. A well-behaved program must test an optional data member before using its value:

Python

```
try:
    ...
except ex:
    if ex.optionalMember:
        print("optionalMember = " + str(ex.optionalMember))
    else:
        print("optionalMember is unset")
```

The `Ice.Unset` marker value has different semantics than `None`. Since `None` is a legal value for certain Slice types, the Ice run time requires a separate marker value so that it can determine whether an optional value is set. An optional value set to `None` is considered to be set. If you need to distinguish between an unset value and a value set to `None`, you can do so as follows:

Python

```
try:
    ...
except ex:
    if ex.optionalMember is Ice.Unset:
        print("optionalMember is unset")
    elif ex.optionalMember is None:
        print("optionalMember is None")
    else:
        print("optionalMember = " + str(ex.optionalMember))
```


Copyright 2016, ZeroC, Inc.

Although Python provides no standard mechanism for restricting access to an object's attributes, by convention an attribute whose name begins with an underscore signals the author's intent that the attribute should only be accessed by the class itself or by one of its subclasses. You can employ this convention in your Slice classes using the `protected` metadata directive. The presence of this directive causes the Slice compiler to prepend an underscore to the mapped name of the data member. For example, the `TimeOfDay` class shown below has the `protected` metadata directive applied to each of its data members:

Slice

```
class TimeOfDay {
    ["protected"] short hour;   // 0 - 23
    ["protected"] short minute; // 0 - 59
    ["protected"] short second; // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Slice compiler produces the following generated code for this definition:

Python

```
class TimeOfDay(Ice.Object):
    def __init__(self, hour=0, minute=0, second=0):
        # ...
        self._hour = hour
        self._minute = minute
        self._second = second

    # ...

    #
    # Operation signatures.
    #
    # def format(self, current=None):
```

For a class in which all of the data members are protected, the metadata directive can be applied to the class itself rather than to each member individually. For example, we can rewrite the `TimeOfDay` class as follows:

Slice

```
["protected"] class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

Class Constructors in Python

Classes have a constructor that assigns to each data member a default value appropriate for its type:

Copyright 2016, ZeroC, Inc.

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	None
`dictionary`	None
`class/interface`	None

You can also declare different default values for data members of primitive and enumerated types.

For derived classes, the constructor has one parameter for each of the base class's data members, plus one parameter for each of the derived class's data members, in base-to-derived order.

You can invoke this constructor in one of two ways:

- Provide values for all members, including optional members, in the order of declaration:

  Python

  ```
  t = TimeOfDay(12, 33, 45)
  t2 = TimeOfDay(14, 7) # second defaults to 0
  ```

  Pass `Ice.Unset` as the value of any optional member you want to be unset.

- Used named arguments to specify values for certain members and in any order:

  Python

  ```
  t = TimeOfDay(minute=33, hour=12)
  ```

Class Operations in Python

Operations of classes are mapped to methods in the generated class. This means that, if a class contains operations (such as the `format` operation of our `TimeOfDay` class), you must provide an implementation of the operation in a class that is derived from the generated class. For example:

Copyright 2016, ZeroC, Inc.

Python

```
class TimeOfDayI(TimeOfDay):
    def __init__(self, hour=0, minute=0, second=0):
        TimeOfDay.__init__(self, hour, minute, second)

    def format(self, current=None):
        return "%02d:%02d:%02d" % (self.hour, self.minute, self.second)
```

A Slice class such as `TimeOfDay` that declares or inherits an operation is inherently abstract. Python does not support the notion of abstract classes or abstract methods, therefore the mapping merely summarizes the required method signatures in a comment for your convenience. Furthermore, the mapping generates code in the constructor of an abstract class to prevent it from being instantiated directly; any attempt to do so raises a `RuntimeError` exception.

You may notice that the mapping for an operation adds an optional trailing parameter named `current`. For now, you can ignore this parameter and pretend it does not exist.

Class Factories in Python

Having created a class such as `TimeOfDayI`, we have an implementation and we can instantiate the `TimeOfDayI` class, but we cannot receive it as the return value or as an out-parameter from an operation invocation. To see why, consider the following simple interface:

Slice

```
interface Time {
    TimeOfDay get();
};
```

When a client invokes the `get` operation, the Ice run time must instantiate and return an instance of the `TimeOfDay` class. However, `TimeOfDay` is an abstract class that cannot be instantiated. Unless we tell it, the Ice run time cannot magically know that we have created a `TimeOfDayI` class that implements the abstract `format` operation of the `TimeOfDay` abstract class. In other words, we must provide the Ice run time with a factory that knows that the `TimeOfDay` abstract class has a `TimeOfDayI` concrete implementation. The `Ice::Communicator` interface provides us with the necessary operations:

Slice

```
module Ice {
    local interface ObjectFactory {
        Object create(string type);
        void destroy();
    };

    local interface Communicator {
        void addObjectFactory(ObjectFactory factory, string id);
        ObjectFactory findObjectFactory(string id);
        // ...
    };
};
```


Copyright 2016, ZeroC, Inc.

To supply the Ice run time with a factory for our `TimeOfDayI` class, we must implement the `ObjectFactory` interface:

Python

```
class ObjectFactory(Ice.ObjectFactory):
    def create(self, type):
        if type == M.TimeOfDay.ice_staticId():
            return TimeOfDayI()
        assert(False)
        return None

    def destroy(self):
        # Nothing to do
        pass
```

The object factory's `create` method is called by the Ice run time when it needs to instantiate a `TimeOfDay` class. The factory's `destroy` m ethod is called by the Ice run time when its communicator is destroyed.

The `create` method is passed the type ID of the class to instantiate. For our `TimeOfDay` class, the type ID is `"::M::TimeOfDay"`. Our implementation of `create` checks the type ID: if it matches, the method instantiates and returns a `TimeOfDayI` object. For other type IDs, the method asserts because it does not know how to instantiate other types of objects.

Note that we used the `ice_staticId` method to obtain the type ID rather than embedding a literal string. Using a literal type ID string in your code is discouraged because it can lead to errors that are only detected at run time. For example, if a Slice class or one of its enclosing modules is renamed and the literal string is not changed accordingly, a receiver will fail to unmarshal the object and the Ice run time will raise `NoObjectFactoryException`. By using `ice_staticId` instead, we avoid any risk of a misspelled or obsolete type ID, and we can discover earlier whether a Slice class or module has been renamed.

Given a factory implementation, such as our `ObjectFactory`, we must inform the Ice run time of the existence of the factory:

Python

```
ic = ...   # Get Communicator...
ic.addObjectFactory(ObjectFactory(), M.TimeOfDay.ice_staticId())
```

Now, whenever the Ice run time needs to instantiate a class with the type ID `"::M::TimeOfDay"`, it calls the `create` method of the registered `ObjectFactory` instance.

The `destroy` operation of the object factory is invoked by the Ice run time when the communicator is destroyed. This gives you a chance to clean up any resources that may be used by your factory. Do not call `destroy` on the factory while it is registered with the communicator — if you do, the Ice run time has no idea that this has happened and, depending on what your `destroy` implementation is doing, may cause undefined behavior when the Ice run time tries to next use the factory.

The run time guarantees that `destroy` will be the last call made on the factory, that is, `create` will not be called concurrently with `destroy` , and `create` will not be called once `destroy` has been called. However, calls to `create` can be made concurrently.

Note that you cannot register a factory for the same type ID twice: if you call `addObjectFactory` with a type ID for which a factory is registered, the Ice run time throws an `AlreadyRegisteredException`.

Finally, keep in mind that if a class has only data members, but no operations, you need not create and register an object factory to transmit instances of such a class. Only if a class has operations do you have to define and register an object factory.

See Also

- Classes
- Type IDs
- Optional Data Members

Copyright 2016, ZeroC, Inc.

- Python Mapping for Operations
- The Current Object

Copyright 2016, ZeroC, Inc.

Asynchronous Method Invocation (AMI) in Python

Asynchronous Method Invocation (AMI) is the term used to describe the client-side support for the asynchronous programming model. AMI supports both oneway and twoway requests, but unlike their synchronous counterparts, AMI requests never block the calling thread. When a client issues an AMI request, the Ice run time hands the message off to the local transport buffer or, if the buffer is currently full, queues the request for later delivery. The application can then continue its activities and poll or wait for completion of the invocation, or receive a callback when the invocation completes.

AMI is transparent to the server: there is no way for the server to tell whether a client sent a request synchronously or asynchronously.

On this page:

- Basic Asynchronous API in Python
  - Asynchronous Proxy Methods in Python
  - Asynchronous Exception Semantics in Python
- AsyncResult Class in Python
- Polling for Completion in Python
- Completion Callbacks in Python
- Sharing State Between begin_ and end_ Methods in Python
- Asynchronous Oneway Invocations in Python
- Flow Control in Python
- Asynchronous Batch Requests in Python
- Concurrency Semantics for AMI in Python

Basic Asynchronous API in Python

Consider the following simple Slice definition:

Slice

```
module Demo {
    interface Employees {
        string getName(int number);
    };
};
```

Asynchronous Proxy Methods in Python

Besides the synchronous proxy methods, the Python mapping generates the following asynchronous proxy methods:

Python

```
def begin_getName(self, number, _response=None, _ex=None, _sent=None, _ctx=None)
def end_getName(self, result)
```

As you can see, the single `getName` operation results in `begin_getName` and `end_getName` methods. The `begin_ method` optionally accepts a per-invocation context and callbacks.

- The `begin_getName` method sends (or queues) an invocation of `getName`. This method does not block the calling thread.
- The `end_getName` method collects the result of the asynchronous invocation. If, at the time the calling thread calls `end_getName`, the result is not yet available, the calling thread blocks until the invocation completes. Otherwise, if the invocation completed some time before the call to `end_getName`, the method returns immediately with the result.

A client could call these methods as follows:

Copyright 2016, ZeroC, Inc.

Python

```
e = EmployeePrx.checkedCast(...)
r = e.begin_getName(99)

# Continue to do other things here...

name = e.end_getName(r)
```

Because `begin_getName` does not block, the calling thread can do other things while the operation is in progress.

Note that `begin_getName` returns a value of type `AsyncResult`. This value contains the state that the Ice run time requires to keep track of the asynchronous invocation. You must pass the `AsyncResult` that is returned by the `begin_` method to the corresponding `end_` method.

The `begin_` method has one parameter for each in-parameter of the corresponding Slice operation. The `end_` method accepts the `AsyncResult` object as its only argument and returns the out-parameters using the same semantics as for regular synchronous invocations. For example, consider the following operation:

Slice

```
double op(int inp1, string inp2, out bool outp1, out long outp2);
```

The `begin_op` and `end_op` methods have the following signature:

Python

```
def begin_op(self, inp1, inp2, ...)
def end_op(self, result)
```

The call to `end_op` returns the following tuple:

Python

```
doubleValue, outp1, outp2 = p.end_op(result)
```

Asynchronous Exception Semantics in Python

If an invocation raises an exception, the exception is thrown by the `end_` method, even if the actual error condition for the exception was encountered during the `begin_` method ("on the way out"). The advantage of this behavior is that all exception handling is located with the code that calls the `end_` method (instead of being present twice, once where the `begin_` method is called, and again where the `end_` method is called).

There is one exception to the above rule: if you destroy the communicator and then make an asynchronous invocation, the `begin_` method throws `CommunicatorDestroyedException`. This is necessary because, once the run time is finalized, it can no longer throw an exception from the `end_` method.

Copyright 2016, ZeroC, Inc.

The only other exception that is thrown by the `begin_` and `end_` methods is `RuntimeError`. This exception indicates that you have used the API incorrectly. For example, the `begin_` method throws this exception if you call an operation that has a return value or out-parameters on a oneway proxy. Similarly, the `end_` method throws this exception if you use a different proxy to call the `end_` method than the proxy you used to call the `begin_` method, or if the `AsyncResult` you pass to the `end_` method was obtained by calling the `begin_` method for a different operation.

`AsyncResult` Class in Python

The `AsyncResult` that is returned by the `begin_` method encapsulates the state of the asynchronous invocation:

Python

```
class AsyncResult:
    def cancel()

    def getCommunicator()
    def getConnection()
    def getProxy()
    def getOperation()

    def isCompleted()
    def waitForCompleted()

    def isSent()
    def waitForSent()

    def throwLocalException()

    def sentSynchronously()
```

The methods have the following semantics:

- `cancel()`
  This method prevents a queued invocation from being sent or, if the invocation has already been sent, ignores a reply if the server sends one. `cancel` is a local operation and has no effect on the server. A canceled invocation is considered to be completed, meaning `isCompleted` returns true, and the result of the invocation is an `Ice.InvocationCanceledException`.
- `getCommunicator()`
  This method returns the communicator that sent the invocation.
- `getConnection()`
  This method returns the connection that was used for the invocation. Note that, for typical asynchronous proxy invocations, this method returns a nil value because the possibility of automatic retries means the connection that is currently in use could change unexpectedly. The `getConnection` method only returns a non-nil value when the `AsyncResult` object is obtained by calling `begin_flushBatchRequests` on a `Connection` object.
- `getProxy()`
  This method returns the proxy that was used to call the `begin_` method, or nil if the `AsyncResult` object was not obtained via an asynchronous proxy invocation.
- `getOperation()`
  This method returns the name of the operation.
- `isCompleted()`
  This method returns true if, at the time it is called, the result of an invocation is available, indicating that a call to the `end_` method will not block the caller. Otherwise, if the result is not yet available, the method returns false.
- `waitForCompleted()`

Copyright 2016, ZeroC, Inc.

This method blocks the caller until the result of an invocation becomes available.

- `isSent()`
  When you call the `begin_` method, the Ice run time attempts to write the corresponding request to the client-side transport. If the transport cannot accept the request, the Ice run time queues the request for later transmission. `isSent` returns true if, at the time it is called, the request has been written to the local transport (whether it was initially queued or not). Otherwise, if the request is still queued or an exception occurred before the request could be sent, `isSent` returns false.
- `waitForSent()`
  This method blocks the calling thread until a request has been written to the client-side transport, or an exception occurs. After `waitForSent` returns, `isSent` returns true if the request was successfully written to the client-side transport, or false if an exception occurred. In the case of a failure, you can call the corresponding `end_` method or `throwLocalException` to obtain the exception.
- `throwLocalException()`
  This method throws the local exception that caused the invocation to fail. If no exception has occurred yet, `throwLocalException` does nothing.
- `sentSynchronously()`
  This method returns true if a request was written to the client-side transport without first being queued. If the request was initially queued, `sentSynchronously` returns false (independent of whether the request is still in the queue or has since been written to the client-side transport).

Polling for Completion in Python

The `AsyncResult` methods allow you to poll for call completion. Polling is useful in a variety of cases. As an example, consider the following simple interface to transfer files from client to server:

Slice

```
interface FileTransfer
{
    void send(int offset, ByteSeq bytes);
};
```

The client repeatedly calls `send` to send a chunk of the file, indicating at which offset in the file the chunk belongs. A naïve way to transmit a file would be along the following lines:

Python

```
file = open(...)
ft = FileTransferPrx.checkedCast(...)
chunkSize = ...
offset = 0
while not file.eof():
    bytes = file.read(chunkSize)  # Read a chunk
    ft.send(offset, bytes)        # Send the chunk
    offset += len(bytes.length)
```

This works, but not very well: because the client makes synchronous calls, it writes each chunk on the wire and then waits for the server to receive the data, process it, and return a reply before writing the next chunk. This means that both client and server spend much of their time doing nothing — the client does nothing while the server processes the data, and the server does nothing while it waits for the client to send the next chunk.

Using asynchronous calls, we can improve on this considerably:

Copyright 2016, ZeroC, Inc.

Python

```
file = open(...)
ft = FileTransferPrx.checkedCast(...)
chunkSize = ...
offset = 0

results = []
numRequests = 5

while not file.eof():
    bytes = file.read(chunkSize) # Read a chunk

    # Send up to numRequests + 1 chunks asynchronously.
    r = ft.begin_send(offset, bytes)
    offset += len(bytes)

    # Wait until this request has been passed to the transport.
    r.waitForSent()
    results.append(r)

    # Once there are more than numRequests, wait for the least
    # recent one to complete.
    while len(results) > numRequests:
        r = results[0]
        del results[0]
        r.waitForCompleted()

# Wait for any remaining requests to complete.
while len(results) > 0:
    r = results[0]
    del results[0]
    r.waitForCompleted()
```

With this code, the client sends up to `numRequests + 1` chunks before it waits for the least recent one of these requests to complete. In other words, the client sends the next request without waiting for the preceding request to complete, up to the limit set by `numRequests`. In effect, this allows the client to "keep the pipe to the server full of data": the client keeps sending data, so both client and server continuously do work.

Obviously, the correct chunk size and value of `numRequests` depend on the bandwidth of the network as well as the amount of time taken by the server to process each request. However, with a little testing, you can quickly zoom in on the point where making the requests larger or queuing more requests no longer improves performance. With this technique, you can realize the full bandwidth of the link to within a percent or two of the theoretical bandwidth limit of a native socket connection.

Completion Callbacks in Python

The `begin_` method accepts three optional callback arguments that allow you to be notified asynchronously when a request completes. Here are the corresponding methods for the `getName` operation:

Copyright 2016, ZeroC, Inc.

Python

```
def begin_getName(self, number, _response=None, _ex=None, _sent=None, _ctx=None)
```

The value you pass for the response callback (`_response`), the exception callback (`_ex`), or the sent callback (`_sent`) argument must be a *callable object* such as a function or method. The response callback is invoked when the request completes successfully, and the exception callback is invoked when the operation raises an exception. (The sent callback is primarily used for flow control.)

For example, consider the following callbacks for an invocation of the `getName` operation:

Python

```
def getNameCB(name):
    print "Name is: " + name

def failureCB(ex):
    print "Exception is: " + str(ex)
```

The response callback parameters depend on the operation signature. If the operation has a non-`void` return type, the first parameter of the response callback is the return value. The return value (if any) is followed by a parameter for each out-parameter of the corresponding Slice operation, in the order of declaration.

The exception callback is invoked if the invocation fails because of an Ice run time exception, or if the operation raises a user exception.

To inform the Ice run time that you want to receive callbacks for the completion of the asynchronous call, you pass the callbacks to the `begin_` method:

Python

```
e = EmployeesPrx.checkedCast(...)

e.begin_getName(99, getNameCB, failureCB)
```

Although the signature of an asynchronous proxy method implies that all of the callbacks are optional and therefore can be supplied in any combination, Ice enforces the following semantics at run time:

- If you omit all callbacks, you must call the `end_` method explicitly as described earlier.
- If you supply either a response callback or a sent callback (or both), you must also supply an exception callback.
- You may omit the response callback for an operation that returns no data (that is, an operation with a `void` return type and no out-parameters).

Sharing State Between `begin_` and `end_` Methods in Python

It is common for the `end_` method to require access to some state that is established by the code that calls the `begin_` method. As an example, consider an application that asynchronously starts a number of operations and, as each operation completes, needs to update different user interface elements with the results. In this case, the `begin_` method knows which user interface element should receive the update, and the `end_` method needs access to that element.

Assuming that we have a `Widget` class that designates a particular user interface element, you could pass different widgets by storing the widget to be used as a member of a callback class:

Copyright 2016, ZeroC, Inc.

Python

```
class MyCallback(object):
    def __init__(self, w):
        self._w = w

    def getNameCB(self, name):
        self._w.writeString(name)

    def failureCB(self, ex):
        print "Exception is: " + str(ex)
```

For this example, we assume that widgets have a `writeString` method that updates the relevant UI element.

When you call the `begin_` method, you pass the appropriate callback instance to inform the `end_` method how to update the display:

Python

```
e = EmployeesPrx.checkedCast(...)
widget1 = ...
widget2 = ...

# Invoke the getName operation with different widget callbacks.
cb1 = MyCallback(widget1)
e.begin_getName(99, cb1.getNameCB, cb1.failureCB)
cb2 = MyCallback(widget2)
e.begin_getName(24, cb2.getNameCB, cb2.failureCB)
```

The callback class provides a simple and effective way for you to pass state between the point where an operation is invoked and the point where its results are processed. Moreover, if you have a number of operations that share common state, you can pass the same callback instance to multiple invocations. (If you do this, your callback methods may need to use synchronization.)

For those situations in which a stateless callback is preferred, you can use a lambda function to pass state to a callback. Consider the following example:

Copyright 2016, ZeroC, Inc.

Python

```
def getNameCB(name, w):
    w.writeString(name)

def failureCB(ex):
    print "Exception is: " + str(ex)

e = EmployeesPrx.checkedCast(...)
widget1 = ...
widget2 = ...

# Use lambda functions to pass state.
e.begin_getName(99, lambda name: getNameCB(name, widget1), failureCB)
e.begin_getName(24, lambda name: getNameCB(name, widget2), failureCB)
```

This strategy eliminates the need to encapsulate shared state in a callback class. Since lambda functions can refer to variables in the enclosing scope, they provide a convenient way to pass state directly to your callback.

Asynchronous Oneway Invocations in Python

You can invoke operations via oneway proxies asynchronously, provided the operation has `void` return type, does not have any out-parameters, and does not raise user exceptions. If you call the `begin_` method on a oneway proxy for an operation that returns values or raises a user exception, the `begin_` method throws a `RuntimeError`.

The callback signatures look exactly as for a twoway invocation, but the response method is never called and may be omitted.

Flow Control in Python

Asynchronous method invocations never block the thread that calls the `begin_` method: the Ice run time checks to see whether it can write the request to the local transport. If it can, it does so immediately in the caller's thread. (In that case, `AsyncResult.sentSynchronously` returns true.) Alternatively, if the local transport does not have sufficient buffer space to accept the request, the Ice run time queues the request internally for later transmission in the background. (In that case, `AsyncResult.sentSynchronously` returns false.)

This creates a potential problem: if a client sends many asynchronous requests at the time the server is too busy to keep up with them, the requests pile up in the client-side run time until, eventually, the client runs out of memory.

The API provides a way for you to implement flow control by counting the number of requests that are queued so, if that number exceeds some threshold, the client stops invoking more operations until some of the queued operations have drained out of the local transport.

You can supply a sent callback to be notified when the request was successfully sent:

Copyright 2016, ZeroC, Inc.

Python

```
def response(name):
    # ...

def exception(ex):
    # ...

def sent(sentSynchronously):
    # ...
```

You inform the Ice run time that you want to be informed when a call has been passed to the local transport as usual:

Python

```
e.begin_getName(99, response, exception, sent)
```

If the Ice run time can immediately pass the request to the local transport, it does so and invokes the sent callback from the thread that calls the `begin_` method. On the other hand, if the run time has to queue the request, it calls the sent callback from a different thread once it has written the request to the local transport. The boolean `sentSynchronously` parameter indicates whether the request was sent synchronously or was queued.

The sent callback allows you to limit the number of queued requests by counting the number of requests that are queued and decrementing the count when the Ice run time passes a request to the local transport.

Asynchronous Batch Requests in Python

Applications that send batched requests can either flush a batch explicitly or allow the Ice run time to flush automatically. The proxy method `ice_flushBatchRequests` performs an immediate flush using the synchronous invocation model and may block the calling thread until the entire message can be sent. Ice also provides asynchronous versions of this method so you can flush batch requests asynchronously.

`begin_ice_flushBatchRequests` and `end_ice_flushBatchRequests` are proxy methods that flush any batch requests queued by that proxy.

In addition, similar methods are available on the communicator and the `Connection` object that is returned by `AsyncResult.getConnection`. These methods flush batch requests sent via the same communicator and via the same connection, respectively.

Concurrency Semantics for AMI in Python

The Ice run time always invokes your callback methods from a separate thread, with one exception: it calls the sent callback from the thread calling the `begin_` method if the request could be sent synchronously. In the `sent` callback, you know which thread is calling the callback by looking at the `sentSynchronously` parameter.

See Also

- Python Mapping for Operations
- Request Contexts
- Batched Invocations

Copyright 2016, ZeroC, Inc.

Code Generation in Python

The Python mapping supports two forms of code generation: dynamic and static.

On this page:

- Dynamic Code Generation in Python
  - Ice.loadSlice Options in Python
  - Locating Slice Files in Python
  - Loading Multiple Slice Files in Python
- Static Code Generation in Python
  - Compiler Output in Python
  - Include Files in Python
- Static Versus Dynamic Code Generation in Python
  - Application Considerations for Code Generation in Python
  - Mixing Static and Dynamic Code Generation in Python
- slice2py Command-Line Options
- Generating Packages in Python

Dynamic Code Generation in Python

Using dynamic code generation, Slice files are "loaded" at run time and dynamically translated into Python code, which is immediately compiled and available for use by the application. This is accomplished using the `Ice.loadSlice` function, as shown in the following example:

Python

```
Ice.loadSlice("Color.ice")
import M

print "My favorite color is", M.Color.blue
```

For this example, we assume that `Color.ice` contains the following definitions:

Slice

```
module M {
    enum Color { red, green, blue };
};
```

The code imports module `M` after the Slice file is loaded because module `M` is not defined until the Slice definitions have been translated into Python.

`Ice.loadSlice` Options in Python

The `Ice.loadSlice` function behaves like a Slice compiler in that it accepts command-line arguments for specifying preprocessor options and controlling code generation. The arguments must include at least one Slice file.

The function has the following Python definition:

Copyright 2016, ZeroC, Inc.

Python

```
def Ice.loadSlice(cmd, args=[])
```

The command-line arguments can be specified entirely in the first argument, `cmd`, which must be a string. The optional second argument can be used to pass additional command-line arguments as a list; this is useful when the caller already has the arguments in list form. The function always returns `None`.

For example, the following calls to `Ice.loadSlice` are functionally equivalent:

Python

```
Ice.loadSlice("-I/opt/IcePy/slice Color.ice")
Ice.loadSlice("-I/opt/IcePy/slice", ["Color.ice"])
Ice.loadSlice("", ["-I/opt/IcePy/slice", "Color.ice"])
```

In addition to the standard compiler options, `Ice.loadSlice` also supports the following command-line options:

- `--all`
  Generate code for all Slice definitions, including those from included files.
- `--checksum`
  Generate checksums for Slice definitions.

Locating Slice Files in Python

If your Slice files depend on Ice types, you can avoid hard-coding the path name of your Ice installation directory by calling the `Ice.getSliceDir` function:

Python

```
Ice.loadSlice("-I" + Ice.getSliceDir() + " Color.ice")
```

This function attempts to locate the `slice` subdirectory of your Ice installation using an algorithm that succeeds for the following scenarios:

- Installation of a binary Ice archive
- Installation of an Ice source distribution using `make install`
- Installation via a Windows installer
- RPM installation on Linux
- Execution inside a compiled Ice source distribution

If the `slice` subdirectory can be found, `getSliceDir` returns its absolute path name, otherwise the function returns `None`.

Loading Multiple Slice Files in Python

You can specify as many Slice files as necessary in a single invocation of `Ice.loadSlice`, as shown below:

Copyright 2016, ZeroC, Inc.

Python

```
Ice.loadSlice("Syscall.ice Process.ice")
```

Alternatively, you can call `Ice.loadSlice` several times:

Python

```
Ice.loadSlice("Syscall.ice")
Ice.loadSlice("Process.ice")
```

If a Slice file includes another file, the default behavior of `Ice.loadSlice` generates Python code only for the named file. For example, suppose `Syscall.ice` includes `Process.ice` as follows:

Slice

```
// Syscall.ice
#include <Process.ice>
...
```

If you call `Ice.loadSlice("-I. Syscall.ice")`, Python code is not generated for the Slice definitions in `Process.ice` or for any definitions that may be included by `Process.ice`. If you also need code to be generated for included files, one solution is to load them individually in subsequent calls to `Ice.loadSlice`. However, it is much simpler, not to mention more efficient, to use the `--all` option instead:

Python

```
Ice.loadSlice("--all -I. Syscall.ice")
```

When you specify `--all`, `Ice.loadSlice` generates Python code for all Slice definitions included directly or indirectly from the named Slice files.

There is no harm in loading a Slice file multiple times, aside from the additional overhead associated with code generation. For example, this situation could arise when you need to load multiple top-level Slice files that happen to include a common subset of nested files. Suppose that we need to load both `Syscall.ice` and `Kernel.ice`, both of which include `Process.ice`. The simplest way to load both files is with a single call to `Ice.loadSlice`:

Python

```
Ice.loadSlice("--all -I. Syscall.ice Kernel.ice")
```

Although this invocation causes the Ice extension to generate code twice for `Process.ice`, the generated code is structured so that the

Copyright 2016, ZeroC, Inc.

interpreter ignores duplicate definitions. We could have avoided generating unnecessary code with the following sequence of steps:

Python

```
Ice.loadSlice("--all -I. Syscall.ice")
Ice.loadSlice("-I. Kernel.ice")
```

In more complex cases, however, it can be difficult or impossible to completely avoid this situation, and the overhead of code generation is usually not significant enough to justify such an effort.

Static Code Generation in Python

You should be familiar with static code generation if you have used other Slice language mappings, such as C++ or Java. Using static code generation, the Slice compiler `slice2py` generates Python code from your Slice definitions.

Compiler Output in Python

For each Slice file `X.ice`, `slice2py` generates Python code into a file named `X_ice.py` in the output directory.

Using the file name `X.py` would create problems if `X.ice` defined a module named `X`, therefore the suffix `_ice` is appended to the name of the generated file.

The default output directory is the current working directory, but a different directory can be specified using the `--output-dir` option.

In addition to the generated file, `slice2py` creates a Python package for each Slice module it encounters. A Python package is nothing more than a subdirectory that contains a file with a special name (`__init__.py`). This file is executed automatically by Python when a program first imports the package. It is created by `slice2py` and must not be edited manually. Inside the file is Python code to import the generated files that contain definitions in the Slice module of interest.

For example, the Slice files `Process.ice` and `Syscall.ice` both define types in the Slice module `OS`. First we present `Process.ice`:

Slice

```
module OS {
    interface Process {
        void kill();
    };
};
```

And here is `Syscall.ice`:

Slice

```
#include <Process.ice>
module OS {
    interface Syscall {
        Process getProcess(int pid);
    };
};
```


Copyright 2016, ZeroC, Inc.

Next, we translate these files using the Slice compiler:

```
> slice2py -I. Process.ice Syscall.ice
```

If we list the contents of the output directory, we see the following entries:

Python

```
OS/
Process_ice.py
Syscall_ice.py
```

The subdirectory `OS` is the Python package that `slice2py` created for the Slice module `OS`. Inside this directory is the special file `__init__.py` that contains the following statements:

Python

```
import Process_ice
import Syscall_ice
```

Now when a Python program executes `import OS`, the two files `Process_ice.py` and `Syscall_ice.py` are implicitly imported.

Subsequent invocations of `slice2py` for Slice files that also contain definitions in the `OS` module result in additional `import` statements being added to `OS/__init__.py`. Be aware, however, that `import` statements may persist in `__init__.py` files after a Slice file is renamed or becomes obsolete. This situation may manifest itself as a run-time error if the interpreter can no longer locate the generated file while attempting to import the package. It may also cause more subtle problems, if an obsolete generated file is still present and being loaded unintentionally. In general, it is advisable to remove the package directory and regenerate it whenever the set of Slice files changes.

A Python program may also import a generated file explicitly, using a statement such as `import Process_ice`. Typically, however, it is more convenient to import the Python module once, rather than importing potentially several individual files that comprise the module, especially when you consider that the program must still import the module explicitly in order to make its definitions available. For example, it is much simpler to state

Python

```
import OS
```

rather than the following alternative:

Python

```
import Process_ice
import Syscall_ice
import OS
```


Copyright 2016, ZeroC, Inc.

Include Files in Python

It is important to understand how `slice2py` handles include files. In the absence of the `--all` option, the compiler does not generate Python code for Slice definitions in included files. Rather, the compiler translates Slice `#include` statements into Python `import` statements in the following manner:

1. Determine the full pathname of the include file.
2. Create the shortest possible relative pathname for the include file by iterating over each of the include directories (specified using the `-I` option) and removing the leading directory from the include file if possible.
   For example, if the full pathname of an include file is `/opt/App/slice/OS/Process.ice`, and we specified the options `-I/opt/App` and `-I/opt/App/slice`, then the shortest relative pathname is `OS/Process.ice` after removing `/opt/App/slice`.
3. Replace any slashes with underscores, remove the `.ice` extension, and append `_ice`. Continuing our example from the previous step, the translated import statement becomes
   `import OS_Process_ice`

There is a potential problem here that must be addressed. The generated `import` statement shown above expects to find the file `OS_Process_ice.py` somewhere in Python's search path. However, `slice2py` uses a different default name, `Process_ice.py`, when it compiles `Process.ice`. To resolve this issue, we must use the `--prefix` option when compiling `Process.ice`:

```
> slice2py --prefix OS_ Process.ice
```

The `--prefix` option causes the compiler to prepend the specified prefix to the name of each generated file. When executed, the above command creates the desired file name: `OS_Process_ice.py`.

It should be apparent by now that generating Python code for a complex Ice application requires a bit of planning. In particular, it is imperative that you be consistent in your use of `#include` statements, include directories, and `--prefix` options to ensure that the correct file names are used at all times.

Of course, these precautionary steps are only necessary when you are compiling Slice files individually. An alternative is to use the `--all` option and generate Python code for all of your Slice definitions into one Python source file. If you do not have a suitable Slice file that includes all necessary Slice definitions, you could write a "master" Slice file specifically for this purpose.

Static Versus Dynamic Code Generation in Python

There are several issues to consider when evaluating your requirements for code generation.

Application Considerations for Code Generation in Python

The requirements of your application generally dictate whether you should use dynamic or static code generation. Dynamic code generation is convenient for a number of reasons:

- it avoids the intermediate compilation step required by static code generation
- it makes the application more compact because the application requires only the Slice files, not the assortment of files and directories produced by static code generation
- it reduces complexity, which is especially helpful during testing, or when writing short or transient programs.

Static code generation, on the other hand, is appropriate in many situations:

- when an application uses a large number of Slice definitions and the startup delay must be minimized
- when it is not feasible to deploy Slice files with the application
- when a number of applications share the same Slice files
- when Python code is required in order to utilize third-party Python tools.

Mixing Static and Dynamic Code Generation in Python

Using a combination of static and dynamic translation in an application can produce unexpected results. For example, consider a situation where a dynamically-translated Slice file includes another Slice file that was statically translated:

Copyright 2016, ZeroC, Inc.

Slice

```
// Slice
#include <Glacier2/Session.ice>

module App {
    interface SessionFactory {
        Glacier2::Session* createSession();
    };
};
```

The Slice file `Session.ice` is statically translated, as are all of the Slice files included with the Ice run time.

Assuming the above definitions are saved in `App.ice`, let's execute a simple Python script:

Python

```
# Python
import Ice
Ice.loadSlice("-I/opt/Ice/slice App.ice")

import Glacier2
class MyVerifier(Glacier2.PermissionsVerifier): # Error
    def checkPermissions(self, userId, password):
        return (True, "")
```

The code looks reasonable, but running it produces the following error:

```
'module' object has no attribute 'PermissionsVerifier'
```

Normally, importing the Glacier2 module as we have done here would load all of the Python code generated for the Glacier2 Slice files. However, since `App.ice` has already included a subset of the Glacier2 definitions, the Python interpreter ignores any subsequent requests to import the entire module, and therefore the `PermissionsVerifier` type is not present.

One way to address this problem is to import the statically-translated modules first, prior to loading Slice files dynamically:

Python

```
# Python
import Ice, Glacier2 # Import Glacier2 before App.ice is loaded
Ice.loadSlice("-I/opt/Ice/slice App.ice")

class MyVerifier(Glacier2.PermissionsVerifier): # OK
    def checkPermissions(self, userId, password):
        return (True, "")
```


Copyright 2016, ZeroC, Inc.

The disadvantage of this approach in a non-trivial application is that it breaks encapsulation, forcing one Python module to know what other modules are doing. For example, suppose we place our `PermissionsVerifier` implementation in a module named `verifier.py`:

Python

```
# Python
import Glacier2
class MyVerifier(Glacier2.PermissionsVerifier):
    def checkPermissions(self, userId, password):
        return (True, "")
```

Now that the use of Glacier2 definitions is encapsulated in `verifier.py`, we would like to remove references to Glacier2 from the main script:

Python

```
# Python
import Ice
Ice.loadSlice("-I/opt/Ice/slice App.ice")
...
import verifier # Error
v = verifier.MyVerifier()
```

Unfortunately, executing this script produces the same error as before. To fix it, we have to break the `verifier` module's encapsulation and import the `Glacier2` module in the main script because we know that the `verifier` module requires it:

Python

```
# Python
import Ice, Glacier2
Ice.loadSlice("-I/opt/Ice/slice App.ice")
...
import verifier # OK
v = verifier.MyVerifier()
```

Although breaking encapsulation in this way might offend our sense of good design, it is a relatively minor issue.

Another solution is to import the necessary submodules explicitly. We can safely remove the Glacier2 reference from our main script after rewriting `verifier.py` as shown below:

Copyright 2016, ZeroC, Inc.

Python

```
# Python
import Glacier2_PermissionsVerifier_ice
import Glacier2
class MyVerifier(Glacier2.PermissionsVerifier):
    def checkPermissions(self, userId, password):
        return (True, "")
```

Using the rules defined for static code generation, we can derive the name of the module containing the code generated for `PermissionsVerifier.ice` and import it directly. We need a second `import` statement to make the Glacier2 definitions accessible in this module.

`slice2py` Command-Line Options

The Slice-to-Python compiler, `slice2py`, offers the following command-line options in addition to the standard options:

- `--all`
  Generate code for all Slice definitions, including those from included files.
- `--checksum`
  Generate checksums for Slice definitions.
- `--prefix PREFIX`
  Use *`PREFIX`* as the prefix for generated file names.

Generating Packages in Python

By default, the scope of a Slice definition determines the module of its mapped Python construct. There are times, however, when applications require greater control over the packaging of generated Python code. For example, consider the following Slice definitions:

Slice

```
module sys {
    interface Process {
        // ...
    };
};
```

Other language mappings can use these Slice definitions as shown, but they present a problem for the Python mapping: the top-level Slice module `sys` conflicts with Python's predefined module sys. A Python application executing the statement `import sys` would import whichever module the interpreter happens to locate first in its search path.

A workaround for this problem is to modify the Slice definitions so that the top-level module no longer conflicts with a predefined Python module, but that may not be feasible in certain situations. For example, the application may already be deployed using other language mappings, in which case the impact of modifying the Slice definitions could represent an unacceptable expense.

The Python mapping could have addressed this issue by considering the names of predefined modules to be reserved, in which case the Slice module `sys` would be mapped to the Python module `_sys`. However, the likelihood of a name conflict is relatively low to justify such a solution, therefore the mapping supports a different approach: global metadata can be used to enclose generated code in a Python package. Our modified Slice definitions demonstrate this feature:

Copyright 2016, ZeroC, Inc.

Slice

```
[["python:package:zeroc"]]
module sys {
    interface Process {
        // ...
    };
};
```

The global metadata directive `python:package:zeroc` causes the mapping to generate all of the code resulting from definitions in this Slice file into the Python package `zeroc`. The net effect is the same as if we had enclosed our Slice definitions in the module `zeroc`: the Slice module `sys` is mapped to the Python module `zeroc.sys`. However, by using metadata we have not affected the semantics of the Slice definitions, nor have we affected other language mappings.

See Also

- Using the Slice Compilers
- Python Mapping for Modules
- Using Slice Checksums in Python
- Metadata

Copyright 2016, ZeroC, Inc.

Using Slice Checksums in Python

The Slice compilers can optionally generate checksums of Slice definitions. For `slice2py`, the `--checksum` option causes the compiler to generate code that adds checksums to the dictionary `Ice.sliceChecksums`. The checksums are installed automatically when the Python code is first imported; no action is required by the application.

In order to verify a server's checksums, a client could simply compare the dictionaries using the comparison operator. However, this is not feasible if it is possible that the server might return a superset of the client's checksums. A more general solution is to iterate over the local checksums as demonstrated below:

Python

```
serverChecksums = ...
for i in Ice.sliceChecksums:
    if not serverChecksums.has_key(i):
        # No match found for type id!
    elif Ice.sliceChecksums[i] != serverChecksums[i]:
        # Checksum mismatch!
```

In this example, the client first verifies that the server's dictionary contains an entry for each Slice type ID, and then it proceeds to compare the checksums.

See Also

- Slice Checksums

Copyright 2016, ZeroC, Inc.

Example of a File System Client in Python

This page presents a very simple client to access a server that implements the file system we developed in Slice for a Simple File System. The Python code shown here hardly differs from the code you would write for an ordinary Python program. This is one of the biggest advantages of using Ice: accessing a remote object is as easy as accessing an ordinary, local Python object. This allows you to put your effort where you should, namely, into developing your application logic instead of having to struggle with arcane networking APIs. This is true for the server side as well, meaning that you can develop distributed applications easily and efficiently.

We now have seen enough of the client-side Python mapping to develop a complete client to access our remote file system. For reference, here is the Slice definition once more:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };

    exception GenericError {
        string reason;
    };

    sequence<string> Lines;

    interface File extends Node {
        idempotent Lines read();
        idempotent void write(Lines text) throws GenericError;
    };

    sequence<Node*> NodeSeq;

    interface Directory extends Node {
        idempotent NodeSeq list();
    };
};
```

To exercise the file system, the client does a recursive listing of the file system, starting at the root directory. For each node in the file system, the client shows the name of the node and whether that node is a file or directory. If the node is a file, the client retrieves the contents of the file and prints them.

The body of the client code looks as follows:

Python

```
import sys, traceback, Ice, Filesystem

# Recursively print the contents of directory "dir"
# in tree fashion. For files, show the contents of
# each file. The "depth" parameter is the current
# nesting level (for indentation).

def listRecursive(dir, depth):
    indent = ''
    depth = depth + 1
```


Copyright 2016, ZeroC, Inc.

```
    for i in range(depth):
        indent = indent + '\t'

    contents = dir.list()

    for node in contents:
        subdir = Filesystem.DirectoryPrx.checkedCast(node)
        file = Filesystem.FilePrx.uncheckedCast(node)
        print indent + node.name(),
        if subdir:
            print "(directory):"
            listRecursive(subdir, depth)
        else:
            print "(file):"
            text = file.read()
            for line in text:
                print indent + "\t" + line

status = 0
ic = None
try:
    # Create a communicator
    #
    ic = Ice.initialize(sys.argv)

    # Create a proxy for the root directory
    #
    obj = ic.stringToProxy("RootDir:default -p 10000")

    # Down-cast the proxy to a Directory proxy
    #
    rootDir = Filesystem.DirectoryPrx.checkedCast(obj)

    # Recursively list the contents of the root directory
    #
    print "Contents of root directory:"
    listRecursive(rootDir, 0)
except:
    traceback.print_exc()
    status = 1

if ic:
    # Clean up
    #
    try:
        ic.destroy()
    except:
        traceback.print_exc()
        status = 1
```


Copyright 2016, ZeroC, Inc.

```
sys.exit(status)
```

The program first defines the `listRecursive` function, which is a helper function to print the contents of the file system, and the main program follows. Let us look at the main program first:

1. The structure of the code follows what we saw in Hello World Application. After initializing the run time, the client creates a proxy to the root directory of the file system. For this example, we assume that the server runs on the local host and listens using the default protocol (TCP/IP) at port 10000. The object identity of the root directory is known to be `RootDir`.
2. The client down-casts the proxy to `DirectoryPrx` and passes that proxy to `listRecursive`, which prints the contents of the file system.

Most of the work happens in `listRecursive`. The function is passed a proxy to a directory to list, and an indent level. (The indent level increments with each recursive call and allows the code to print the name of each node at an indent level that corresponds to the depth of the tree at that node.) `listRecursive` calls the `list` operation on the directory and iterates over the returned sequence of nodes:

1. The code does a `checkedCast` to narrow the `Node` proxy to a `Directory` proxy, as well as an `uncheckedCast` to narrow the `Node` proxy to a `File` proxy. Exactly one of those casts will succeed, so there is no need to call `checkedCast` twice: if the `Node` *is-a* `Directory`, the code uses the `DirectoryPrx` returned by the `checkedCast`; if the `checkedCast` fails, we *know* that the Node *is-a* File and, therefore, an `uncheckedCast` is sufficient to get a `FilePrx`.
   In general, if you know that a down-cast to a specific type will succeed, it is preferable to use an `uncheckedCast` instead of a `checkedCast` because an `uncheckedCast` does not incur any network traffic.
2. The code prints the name of the file or directory and then, depending on which cast succeeded, prints `"(directory)"` or `"(file)"` following the name.
3. The code checks the type of the node:
   - If it is a directory, the code recurses, incrementing the indent level.
   - If it is a file, the code calls the `read` operation on the file to retrieve the file contents and then iterates over the returned sequence of lines, printing each line.

Assume that we have a small file system consisting of a two files and a directory as follows:

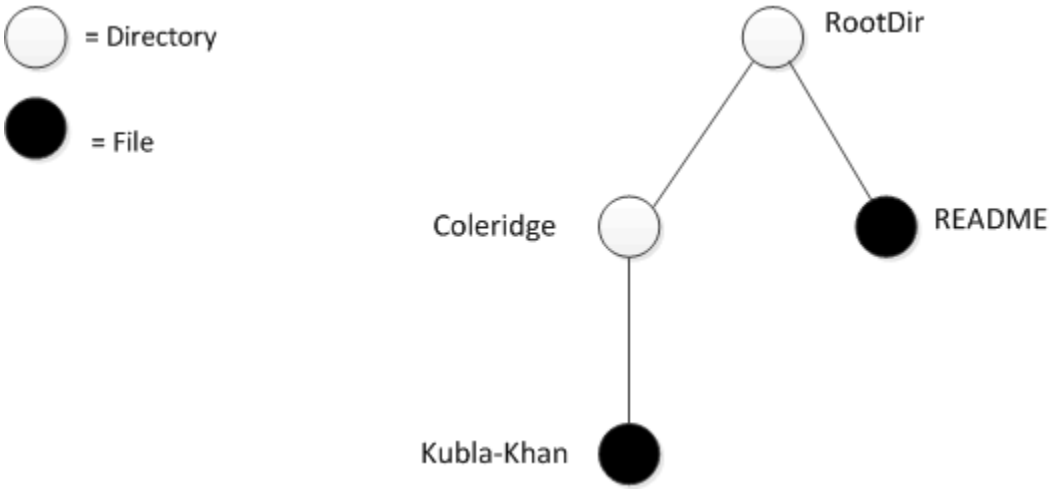

A small file system.

The output produced by the client for this file system is:

Copyright 2016, ZeroC, Inc.

```
Contents of root directory:
        README (file):
                This file system contains a collection of poetry.
        Coleridge (directory):
                Kubla_Khan (file):
                        In Xanadu did Kubla Khan
                        A stately pleasure-dome decree:
                        Where Alph, the sacred river, ran
                        Through caverns measureless to man
                        Down to a sunless sea.
```

Note that, so far, our client (and server) are not very sophisticated:

- The protocol and address information are hard-wired into the code.
- The client makes more remote procedure calls than strictly necessary; with minor redesign of the Slice definitions, many of these calls can be avoided.

We will see how to address these shortcomings in our discussions of IceGrid and object life cycle.

See Also

- Hello World Application
- Slice for a Simple File System
- Example of a File System Server in Python
- Object Life Cycle
- IceGrid

Copyright 2016, ZeroC, Inc.

Server-Side Slice-to-Python Mapping

The mapping for Slice data types to Python is identical on the client side and server side. This means that everything in the Client-Side Slice-to-Python Mapping section also applies to the server side. However, for the server side, there are a few additional things you need to know — specifically, how to:

- Initialize and finalize the server-side run time
- Implement servants
- Pass parameters and throw exceptions
- Create servants and register them with the Ice run time.

Although the examples in this chapter are simple, they accurately reflect the basics of writing an Ice server. Of course, for more sophisticated servers, you will be using additional APIs, for example, to improve performance or scalability. However, these APIs are all described in Slice, so, to use these APIs, you need not learn any Python mapping rules beyond those described here.

The Python interpreter's Global Interpreter Lock (GIL) can impact an Ice server's ability to utilize multiple processing cores. Refer to our FAQ for for more information.

Topics

- The Server-Side main Program in Python
- Server-Side Python Mapping for Interfaces
- Parameter Passing in Python
- Raising Exceptions in Python
- Object Incarnation in Python
- Asynchronous Method Dispatch (AMD) in Python
- Example of a File System Server in Python

Copyright 2016, ZeroC, Inc.

The Server-Side main Program in Python

This section discussed how to initialize and finalize the server-side run time.

On this page:

- Initializing and Finalizing the Server-Side Run Time in Python
- The Ice.Application Class in Python
  - Using Ice.Application on the Client Side in Python
  - Catching Signals in Python
  - Ice.Application and Properties in Python
  - Limitations of Ice.Application in Python

Initializing and Finalizing the Server-Side Run Time in Python

The main entry point to the Ice run time is represented by the local interface `Ice::Communicator`. As for the client side, you must initialize the Ice run time by calling `Ice.initialize` before you can do anything else in your server. `Ice.initialize` returns a reference to an instance of `Ice.Communicator`:

Python

```
import sys, traceback, Ice

status = 0
ic = None
try:
    ic = Ice.initialize(sys.argv)
    # ...
except:
    traceback.print_exc()
    status = 1

# ...
```

`Ice.initialize` accepts the argument list that is passed to the program by the operating system. The function scans the argument list for any command-line options that are relevant to the Ice run time; any such options are removed from the argument list so, when `Ice.initialize` returns, the only options and arguments remaining are those that concern your application. If anything goes wrong during initialization, `initialize` throws an exception.

You can pass a second argument of type `InitializationData` to `Ice.initialize`. `InitializationData` is defined as follows:

Python

```
class InitializationData(object):
    def __init__(self):
        self.properties = None
        self.logger = None
        self.threadHook = None
```

You can pass in an instance of this class to set properties for the communicator, establish a logger, and to establish a thread notification hook.

Before leaving your program, you *must* call `Communicator.destroy`. The `destroy` operation is responsible for finalizing the Ice run time. In particular, `destroy` waits for any operation implementations that are still executing in the server to complete. In addition, `destroy` ensures that any outstanding threads are joined with and reclaims a number of operating system resources, such as file descriptors and memory. Never allow your program to terminate without calling `destroy` first; doing so has undefined behavior.

Copyright 2016, ZeroC, Inc.

The general shape of our server-side program is therefore as follows:

Python

```
import sys, traceback, Ice

status = 0
ic = None
try:
    ic = Ice.initialize(sys.argv)
    # ...
except:
    traceback.print_exc()
    status = 1

if ic:
    try:
        ic.destroy()
    except:
        traceback.print_exc()
        status = 1

sys.exit(status)
```

Note that the code places the call to `Ice.initialize` into a `try` block and takes care to return the correct exit status to the operating system. Also note that an attempt to destroy the communicator is made only if the initialization succeeded.

The `Ice.Application` Class in Python

The preceding program structure is so common that Ice offers a class, `Ice.Application`, that encapsulates all the correct initialization and finalization activities. The synopsis of the class is as follows (with some detail omitted for now):

Python

```
class Application(object):

    def __init__(self, signalPolicy=0):

    def main(self, args, configFile=None, initData=None):

    def run(self, args):

    def appName():
        # ...
    appName = staticmethod(appName)

    def communicator():
        # ...
    communicator = staticmethod(communicator)
```


Copyright 2016, ZeroC, Inc.

The intent of this class is that you specialize `Ice.Application` and implement the abstract `run` method in your derived class. Whatever code you would normally place in your main program goes into `run` instead. Using `Ice.Application`, our program looks as follows:

Python

```
import sys, Ice

class Server(Ice.Application):
    def run(self, args):
        # Server code here...
        return 0

app = Server()
status = app.main(sys.argv)
sys.exit(status)
```

You also can call `main` with an optional file name or an `InitializationData` structure. If you pass a configuration file name to `main`, the property settings in this file are overridden by settings in a file identified by the `ICE_CONFIG` environment variable (if defined). Property settings supplied on the command line take precedence over all other settings.

The `Application.main` function does the following:

1. It installs an exception handler. If your code fails to handle an exception, `Application.main` prints the exception information before returning with a non-zero return value.
2. It initializes (by calling `Ice.initialize`) and finalizes (by calling `Communicator.destroy`) a communicator. You can get access to the communicator for your server by calling the static `communicator` accessor.
3. It scans the argument list for options that are relevant to the Ice run time and removes any such options. The argument list that is passed to your `run` method therefore is free of Ice-related options and only contains options and arguments that are specific to your application.
4. It provides the name of your application via the static `appName` member function. The return value from this call is the first element of the argument vector passed to `Application.main`, so you can get at this name from anywhere in your code by calling `Ice.Application.appName` (which is often necessary for error messages).
5. It installs a signal handler that properly shuts down the communicator.
6. It installs a per-process logger if the application has not already configured one. The per-process logger uses the value of the `Ice.ProgramName` property as a prefix for its messages and sends its output to the standard error channel. An application can also specify an alternate logger.

Using `Ice.Application` ensures that your program properly finalizes the Ice run time, whether your server terminates normally or in response to an exception or signal. We recommend that all your programs use this class; doing so makes your life easier. In addition, `Ice.Application` also provides features for signal handling and configuration that you do not have to implement yourself when you use this class.

Using `Ice.Application` on the Client Side in Python

You can use `Ice.Application` for your clients as well: simply implement a class that derives from `Ice.Application` and place the client code into its `run` method. The advantage of this approach is the same as for the server side: `Ice.Application` ensures that the communicator is destroyed correctly even in the presence of exceptions or signals.

Catching Signals in Python

The simple server we developed in Hello World Application had no way to shut down cleanly: we simply interrupted the server from the command line to force it to exit. Terminating a server in this fashion is unacceptable for many real-life server applications: typically, the server has to perform some cleanup work before terminating, such as flushing database buffers or closing network connections. This is particularly important on receipt of a signal or keyboard interrupt to prevent possible corruption of database files or other persistent data.

To make it easier to deal with signals, `Ice.Application` encapsulates Python's signal handling capabilities, allowing you to cleanly shut down on receipt of a signal:

Copyright 2016, ZeroC, Inc.

Python

```
class Application(object):
    # ...
    def destroyOnInterrupt():
        # ...
    destroyOnInterrupt = classmethod(destroyOnInterrupt)

    def shutdownOnInterrupt():
        # ...
    shutdownOnInterrupt = classmethod(shutdownOnInterrupt)

    def ignoreInterrupt():
        # ...
    ignoreInterrupt = classmethod(ignoreInterrupt)

    def callbackOnInterrupt():
        # ...
    callbackOnInterrupt = classmethod(callbackOnInterrupt)

    def holdInterrupt():
        # ...
    holdInterrupt = classmethod(holdInterrupt)

    def releaseInterrupt():
        # ...
    releaseInterrupt = classmethod(releaseInterrupt)

    def interrupted():
        # ...
    interrupted = classmethod(interrupted)

    def interruptCallback(self, sig):
        # Default implementation does nothing.
        pass
```

The methods behave as follows:

- `destroyOnInterrupt`
  This method installs a signal handler that destroys the communicator if it is interrupted. This is the default behavior.
- `shutdownOnInterrupt`
  This method installs a signal handler that shuts down the communicator if it is interrupted.
- `ignoreInterrupt`
  This method causes signals to be ignored.
- `callbackOnInterrupt`
  This method configures `Ice.Application` to invoke `interruptCallback` when a signal occurs, thereby giving the subclass responsibility for handling the signal.
- `holdInterrupt`
  This method temporarily blocks signal delivery.
- `releaseInterrupt`
  This method restores signal delivery to the previous disposition. Any signal that arrives after `holdInterrupt` was called is

Copyright 2016, ZeroC, Inc.

delivered when you call `releaseInterrupt`.

- `interrupted`
  This method returns `True` if a signal caused the communicator to shut down, `False` otherwise. This allows us to distinguish intentional shutdown from a forced shutdown that was caused by a signal. This is useful, for example, for logging purposes.
- `interruptCallback`
  A subclass overrides this method to respond to signals. The function may be called concurrently with any other thread and must not raise exceptions.

By default, `Ice.Application` behaves as if `destroyOnInterrupt` was invoked, therefore our server program requires no change to ensure that the program terminates cleanly on receipt of a signal. (You can disable the signal-handling functionality of `Ice.Application` by passing `1` to the constructor. In that case, signals retain their default behavior, that is, terminate the process.) However, we add a diagnostic to report the occurrence of a signal, so our program now looks like:

Python

```
import sys, Ice

class MyApplication(Ice.Application):
    def run(self, args):

        # Server code here...

        if self.interrupted():
            print self.appName() + ": terminating"

        return 0

app = MyApplication()
status = app.main(sys.argv)
sys.exit(status)
```

Note that, if your server is interrupted by a signal, the Ice run time waits for all currently executing operations to finish. This means that an operation that updates persistent state cannot be interrupted in the middle of what it was doing and cause partial update problems.

`Ice.Application` and Properties in Python

Apart from the functionality shown in this section, `Ice.Application` also takes care of initializing the Ice run time with property values. Properties allow you to configure the run time in various ways. For example, you can use properties to control things such as the thread pool size or port number for a server. The `main` method of `Ice.Application` accepts an optional second parameter allowing you to specify the name of a configuration file that will be processed during initialization.

Limitations of `Ice.Application` in Python

`Ice.Application` is a singleton class that creates a single communicator. If you are using multiple communicators, you cannot use `Ice.Application`. Instead, you must structure your code as we saw in Hello World Application (taking care to always destroy the communicator).
See Also

- Hello World Application
- Properties and Configuration
- Communicator Initialization
- Logger Facility
- Thread Safety

Copyright 2016, ZeroC, Inc.

Server-Side Python Mapping for Interfaces

The server-side mapping for interfaces provides an up-call API for the Ice run time: by implementing methods in a servant class, you provide the hook that gets the thread of control from the Ice server-side run time into your application code.

On this page:

- Skeleton Classes in Python
- Servant Classes in Python
  - Server-Side Normal and idempotent Operations in Python

Skeleton Classes in Python

On the client side, interfaces map to proxy classes. On the server side, interfaces map to *skeleton classes*. A skeleton is an abstract base class from which you derive your servant class and define a method for each operation on the corresponding interface. For example, consider our Slice definition for the `Node` interface:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };
    // ...
};
```

The Python mapping generates the following definition for this interface:

Python

```
class Node(Ice.Object):
    def __init__(self):
        # ...

    #
    # Operation signatures.
    #
    # def name(self, current=None):
```

The important points to note here are:

- As for the client side, Slice modules are mapped to Python modules with the same name, so the skeleton class definitions are part of the `Filesystem` module.
- The name of the skeleton class is the same as the name of the Slice interface (`Node`).
- The skeleton class contains a comment summarizing the method signature of each operation in the Slice interface.
- The skeleton class is an abstract base class because its constructor prevents direct instantiation of the class.
- The skeleton class inherits from `Ice.Object` (which forms the root of the Ice object hierarchy).

Servant Classes in Python

In order to provide an implementation for an Ice object, you must create a servant class that inherits from the corresponding skeleton class. For example, to create a servant for the `Node` interface, you could write:

Copyright 2016, ZeroC, Inc.

Python

```
import Filesystem

class NodeI(Filesystem.Node):
    def __init__(self, name):
        self._name = name

    def name(self, current=None):
        return self._name
```

By convention, servant classes have the name of their interface with an `I`-suffix, so the servant class for the `Node` interface is called `NodeI`. (This is a convention only: as far as the Ice run time is concerned, you can choose any name you prefer for your servant classes.) Note that `NodeI` extends `Filesystem.Node`, that is, it derives from its skeleton class.

As far as Ice is concerned, the `NodeI` class must implement only a single method: the `name` method that is defined in the `Node` interface. This makes the servant class a concrete class that can be instantiated. You can add other member functions and data members as you see fit to support your implementation. For example, in the preceding definition, we added a `_name` member and a constructor. (Obviously, the constructor initializes the `_name` member and the `name` function returns its value.)

Server-Side Normal and `idempotent` Operations in Python

Whether an operation is an ordinary operation or an `idempotent` operation has no influence on the way the operation is mapped. To illustrate this, consider the following interface:

Slice

```
interface Example {
               void normalOp();
    idempotent void idempotentOp();
};
```

The mapping for this interface is shown below:

Python

```
class Example(Ice.Object):
    # ...

    #
    # Operation signatures.
    #
    # def normalOp(self, current=None):
    # def idempotentOp(self, current=None):
```

Note that the signatures of the methods are unaffected by the `idempotent` qualifier.
See Also

- Slice for a Simple File System

Copyright 2016, ZeroC, Inc.

- Python Mapping for Interfaces
- Parameter Passing in Python
- Raising Exceptions in Python

Copyright 2016, ZeroC, Inc.

Parameter Passing in Python

For each `in` parameter of a Slice operation, the Python mapping generates a corresponding parameter for the method in the skeleton. In addition, every operation has a trailing parameter of type `Ice.Current`. For example, the `name` operation of the `Node` interface has no parameters, but the `name` method in a Python servant has a `current` parameter. We will ignore this parameter for now.

Parameter passing on the server side follows the rules for the client side. An operation returning multiple values returns them in a tuple consisting of a non-`void` return value, if any, followed by the `out` parameters in the order of declaration. An operation returning only one value simply returns the value itself.

An operation returns multiple values when it declares multiple out parameters, or when it declares a non-`void` return type and at least one `out` parameter.

To illustrate these rules, consider the following interface that passes string parameters in all possible directions:

Slice

```
interface Example {
    string op1(string sin);
    void op2(string sin, out string sout);
    string op3(string sin, out string sout);
};
```

The generated skeleton class for this interface looks as follows:

Python

```
class Example(Ice.Object):
    def __init__(self):
        # ...

    #
    # Operation signatures.
    #
    # def op1(self, sin, current=None):
    # def op2(self, sin, current=None):
    # def op3(self, sin, current=None):
```

The signatures of the Python methods are identical because they all accept a single `in` parameter, but their implementations differ in the way they return values. For example, we could implement the operations as follows:

Copyright 2016, ZeroC, Inc.

Python

```
class ExampleI(Example):
    def op1(self, sin, current=None):
        print sin             # In params are initialized
        return "Done"         # Return value

    def op2(self, sin, current=None):
        print sin             # In params are initialized
        return "Hello World!" # Out parameter

    def op3(self, sin, current=None):
        print sin             # In params are initialized
        return ("Done", "Hello World!")
```

Notice that `op1` and `op2` return their string values directly, whereas `op3` returns a tuple consisting of the return value followed by the `out` parameter.

This code is in no way different from what you would normally write if you were to pass strings to and from a function; the fact that remote procedure calls are involved does not impact your code in any way. The same is true for parameters of other types, such as proxies, classes, or dictionaries: the parameter passing conventions follow normal Python rules and do not require special-purpose API calls.

See Also

- Server-Side Python Mapping for Interfaces
- Python Mapping for Operations
- Raising Exceptions in Python
- The Current Object

Copyright 2016, ZeroC, Inc.

Raising Exceptions in Python

To throw an exception from an operation implementation, you simply instantiate the exception, initialize it, and throw it. For example:

Python

```
class FileI(Filesystem.File):
    # ...

    def write(self, text, current=None):
        # Try to write the file contents here...
        # Assume we are out of space...
        if error:
            e = Filesystem.GenericError()
            e.reason = "file too large"
            raise e
```

The mapping for exceptions generates a constructor that accepts values for data members, so we can simplify this example by changing our `raise` statement to the following:

Python

```
class FileI(Filesystem.File):
    # ...

    def write(self, text, current=None):
        # Try to write the file contents here...
        # Assume we are out of space...
        if error:
            raise Filesystem.GenericError("file too large")
```

If you throw an arbitrary Python run-time exception, the Ice run time catches the exception and then returns an `UnknownException` to the client. Similarly, if you throw an "impossible" user exception (a user exception that is not listed in the exception specification of the operation), the client receives an `UnknownUserException`.

If you throw a run-time exception, such as `MemoryLimitException`, the client receives an `UnknownLocalException`. For that reason, you should never throw system exceptions from operation implementations. If you do, all the client will see is an `UnknownLocalException`, which does not tell the client anything useful.

Three run-time exceptions are treated specially and not changed to `UnknownLocalException` when returned to the client: `ObjectNotExistException`, `OperationNotExistException`, and `FacetNotExistException`.

See Also

- Run-Time Exceptions
- Server-Side Python Mapping for Interfaces
- Python Mapping for Exceptions
- Versioning

Copyright 2016, ZeroC, Inc.

Object Incarnation in Python

Having created a servant class such as the rudimentary `NodeI` class, you can instantiate the class to create a concrete servant that can receive invocations from a client. However, merely instantiating a servant class is insufficient to incarnate an object. Specifically, to provide an implementation of an Ice object, you must take the following steps:

1. Instantiate a servant class.
2. Create an identity for the Ice object incarnated by the servant.
3. Inform the Ice run time of the existence of the servant.
4. Pass a proxy for the object to a client so the client can reach it.

On this page:

- Instantiating a Python Servant
- Creating an Identity in Python
- Activating a Python Servant
- UUIDs as Identities in Python
- Creating Proxies in Python
  - Proxies and Servant Activation in Python
  - Direct Proxy Creation in Python

Instantiating a Python Servant

Instantiating a servant means to allocate an instance:

Python

```
servant = NodeI("Fred")
```

This statement creates a new `NodeI` instance and assigns its reference to the variable `servant`.

Creating an Identity in Python

Each Ice object requires an identity. That identity must be unique for all servants using the same object adapter.

The Ice object model assumes that all objects (regardless of their adapter) have a globally unique identity.

An Ice object identity is a structure with the following Slice definition:

Slice

```
module Ice {
    struct Identity {
        string name;
        string category;
    };
    // ...
};
```

The full identity of an object is the combination of both the `name` and `category` fields of the `Identity` structure. For now, we will leave the `category` field as the empty string and simply use the `name` field. (The `category` field is most often used in conjunction with servant locators.)

To create an identity, we simply assign a key that identifies the servant to the `name` field of the `Identity` structure:

Copyright 2016, ZeroC, Inc.

Python

```
id = Ice.Identity()
id.name = "Fred" # Not unique, but good enough for now
```

Note that the mapping for structures allows us to write the following equivalent code:

Python

```
id = Ice.Identity("Fred") # Not unique, but good enough for now
```

Activating a Python Servant

Merely creating a servant instance does nothing: the Ice run time becomes aware of the existence of a servant only once you explicitly tell the object adapter about the servant. To activate a servant, you invoke the `add` operation on the object adapter. Assuming that we have access to the object adapter in the `adapter` variable, we can write:

Python

```
adapter.add(servant, id)
```

Note the two arguments to `add`: the servant and the object identity. Calling `add` on the object adapter adds the servant and the servant's identity to the adapter's servant map and links the proxy for an Ice object to the correct servant instance in the server's memory as follows:

1. The proxy for an Ice object, apart from addressing information, contains the identity of the Ice object. When a client invokes an operation, the object identity is sent with the request to the server.
2. The object adapter receives the request, retrieves the identity, and uses the identity as an index into the servant map.
3. If a servant with that identity is active, the object adapter retrieves the servant from the servant map and dispatches the incoming request into the correct member function on the servant.

Assuming that the object adapter is in the active state, client requests are dispatched to the servant as soon as you call `add`.

UUIDs as Identities in Python

The Ice object model assumes that object identities are globally unique. One way of ensuring that uniqueness is to use UUIDs (Universally Unique Identifiers) as identities. The `Ice.generateUUID` function creates such identities:

Python

```
import Ice
print Ice.generateUUID()
```

When executed, this program prints a unique string such as `5029a22c-e333-4f87-86b1-cd5e0fcce509`. Each call to `generateUUID` creates a string that differs from all previous ones.

You can use a UUID such as this to create object identities. For convenience, the object adapter has an operation `addWithUUID` that

Copyright 2016, ZeroC, Inc.

generates a UUID and adds a servant to the servant map in a single step. Using this operation, we can create an identity and register a servant with that identity in a single step as follows:

Python

```
adapter.addWithUUID(NodeI("Fred"))
```

Creating Proxies in Python

Once we have activated a servant for an Ice object, the server can process incoming client requests for that object. However, clients can only access the object once they hold a proxy for the object. If a client knows the server's address details and the object identity, it can create a proxy from a string, as we saw in our first example in Hello World Application. However, creation of proxies by the client in this manner is usually only done to allow the client access to initial objects for bootstrapping. Once the client has an initial proxy, it typically obtains further proxies by invoking operations.

The object adapter contains all the details that make up the information in a proxy: the addressing and protocol information, and the object identity. The Ice run time offers a number of ways to create proxies. Once created, you can pass a proxy to the client as the return value or as an out-parameter of an operation invocation.

Proxies and Servant Activation in Python

The `add` and `addWithUUID` servant activation operations on the object adapter return a proxy for the corresponding Ice object. This means we can write:

Python

```
proxy = adapter.addWithUUID(NodeI("Fred"))
nodeProxy = Filesystem.NodePrx.uncheckedCast(proxy)

# Pass nodeProxy to client...
```

Here, `addWithUUID` both activates the servant and returns a proxy for the Ice object incarnated by that servant in a single step.

Note that we need to use an `uncheckedCast` here because `addWithUUID` returns a proxy of type `Ice.ObjectPrx`.

Direct Proxy Creation in Python

The object adapter offers an operation to create a proxy for a given identity:

Slice

```
module Ice {
    local interface ObjectAdapter {
        Object* createProxy(Identity id);
        // ...
    };
};
```

Note that `createProxy` creates a proxy for a given identity whether a servant is activated with that identity or not. In other words, proxies have a life cycle that is quite independent from the life cycle of servants:

Copyright 2016, ZeroC, Inc.

Python

```
id = Ice.Identity()
id.name = Ice.generateUUID()
proxy = adapter.createProxy(id)
```

This creates a proxy for an Ice object with the identity returned by `generateUUID`. Obviously, no servant yet exists for that object so, if we return the proxy to a client and the client invokes an operation on the proxy, the client will receive an `ObjectNotExistException`. (We examine these life cycle issues in more detail in Object Life Cycle.)

See Also

- Hello World Application
- Python Mapping for Structures
- Server-Side Python Mapping for Interfaces
- Object Adapter States
- Servant Locators
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

Asynchronous Method Dispatch (AMD) in Python

The number of simultaneous synchronous requests a server is capable of supporting is determined by the number of threads in the server's thread pool. If all of the threads are busy dispatching long-running operations, then no threads are available to process new requests and therefore clients may experience an unacceptable lack of responsiveness.

Asynchronous Method Dispatch (AMD), the server-side equivalent of AMI, addresses this scalability issue. Using AMD, a server can receive a request but then suspend its processing in order to release the dispatch thread as soon as possible. When processing resumes and the results are available, the server sends a response explicitly using a callback object provided by the Ice run time.

AMD is transparent to the client, that is, there is no way for a client to distinguish a request that, in the server, is processed synchronously from a request that is processed asynchronously.

In practical terms, an AMD operation typically queues the request data (i.e., the callback object and operation arguments) for later processing by an application thread (or thread pool). In this way, the server minimizes the use of dispatch threads and becomes capable of efficiently supporting thousands of simultaneous clients.

An alternate use case for AMD is an operation that requires further processing after completing the client's request. In order to minimize the client's delay, the operation returns the results while still in the dispatch thread, and then continues using the dispatch thread for additional work.

On this page:

- Enabling AMD with Metadata in Python
- AMD Mapping in Python
- AMD Exceptions in Python
- AMD Example in Python

Enabling AMD with Metadata in Python

To enable asynchronous dispatch, you must add an `["amd"]` metadata directive to your Slice definitions. The directive applies at the interface and the operation level. If you specify `["amd"]` at the interface level, all operations in that interface use asynchronous dispatch; if you specify `["amd"]` for an individual operation, only that operation uses asynchronous dispatch. In either case, the metadata directive *replaces* synchronous dispatch, that is, a particular operation implementation must use synchronous or asynchronous dispatch and cannot use both.

Consider the following Slice definitions:

Slice

```
["amd"] interface I {
    bool isValid();
    float computeRate();
};

interface J {
    ["amd"] void startProcess();
    int endProcess();
};
```

In this example, both operations of interface `I` use asynchronous dispatch, whereas, for interface `J`, `startProcess` uses asynchronous dispatch and `endProcess` uses synchronous dispatch.

Specifying metadata at the operation level (rather than at the interface or class level) minimizes the amount of generated code and, more importantly, minimizes complexity: although the asynchronous model is more flexible, it is also more complicated to use. It is therefore in your best interest to limit the use of the asynchronous model to those operations that need it, while using the simpler synchronous model for the rest.

AMD Mapping in Python

For each AMD operation, the Python mapping emits a dispatch method with the same name as the operation and the suffix `_async`. This method returns `None`. The first parameter is a reference to a callback object, as described below. The remaining parameters comprise the `in` parameters of the operation, in the order of declaration.

Copyright 2016, ZeroC, Inc.

The callback object defines two methods:

- `def ice_response(self, <params>)`
  The `ice_response` method allows the server to report the successful completion of the operation. If the operation has a non-`void` return type, the first parameter to `ice_response` is the return value. Parameters corresponding to the operation's `out` parameters follow the return value, in the order of declaration.
- `def ice_exception(self, ex)`
  The `ice_exception` method allows the server to report an exception.

Neither `ice_response` nor `ice_exception` throw any exceptions to the caller.

Suppose we have defined the following operation:

Slice

```
interface I {
    ["amd"] int foo(short s, out long l);
};
```

The callback interface generated for operation `foo` is shown below:

Python

```
class ...
    #
    # Operation signatures.
    #
    # def ice_response(self, _result, l)
    # def ice_exception(self, ex)
```

The dispatch method for asynchronous invocation of operation `foo` is generated as follows:

Python

```
def foo_async(self, __cb, s)
```

AMD Exceptions in Python

There are two processing contexts in which the logical implementation of an AMD operation may need to report an exception: the dispatch thread (the thread that receives the invocation), and the response thread (the thread that sends the response).

These are not necessarily two different threads: it is legal to send the response from the dispatch thread.

Although we recommend that the callback object be used to report all exceptions to the client, it is legal for the implementation to raise an exception instead, but only from the dispatch thread.

As you would expect, an exception raised from a response thread cannot be caught by the Ice run time; the application's run time environment determines how such an exception is handled. Therefore, a response thread must ensure that it traps all exceptions and sends the appropriate response using the callback object. Otherwise, if a response thread is terminated by an uncaught exception, the request may never be completed and the client might wait indefinitely for a response.

Copyright 2016, ZeroC, Inc.

Whether raised in a dispatch thread or reported via the callback object, user exceptions are validated and local exceptions may undergo translation.

AMD Example in Python

To demonstrate the use of AMD in Ice, let us define the Slice interface for a simple computational engine:

Slice

```
module Demo {
    sequence<float> Row;
    sequence<Row> Grid;

    exception RangeError {};

    interface Model {
        ["amd"] Grid interpolate(Grid data, float factor)
            throws RangeError;
    };
};
```

Given a two-dimensional grid of floating point values and a factor, the `interpolate` operation returns a new grid of the same size with the values interpolated in some interesting (but unspecified) way.

Our servant class derives from `Demo.Model` and supplies a definition for the `interpolate_async` method that creates a `Job` to hold the callback object and arguments, and adds the `Job` to a queue. The method uses a lock to guard access to the queue:

Python

```
class ModelI(Demo.Model):
    def __init__(self):
        self._mutex = threading.Lock()
        self._jobs = []

    def interpolate_async(self, cb, data, factor, current=None):
        self._mutex.acquire()
        try:
            self._jobs.append(Job(cb, data, factor))
        finally:
            self._mutex.release()
```

After queuing the information, the operation returns control to the Ice run time, making the dispatch thread available to process another request. An application thread removes the next `Job` from the queue and invokes `execute`, which uses `interpolateGrid` (not shown) to perform the computational work:

Copyright 2016, ZeroC, Inc.

Python

```
class Job(object):
    def __init__(self, cb, grid, factor):
        self._cb = cb
        self._grid = grid
        self._factor = factor

    def execute(self):
        if not self.interpolateGrid():
            self._cb.ice_exception(Demo.RangeError())
            return
        self._cb.ice_response(self._grid)

    def interpolateGrid(self):
        # ...
```

If `interpolateGrid` returns `False`, then `ice_exception` is invoked to indicate that a range error has occurred. The `return` statement following the call to `ice_exception` is necessary because `ice_exception` does not throw an exception; it only marshals the exception argument and sends it to the client.

If interpolation was successful, `ice_response` is called to send the modified grid back to the client.
See Also

- User Exceptions
- Run-Time Exceptions
- Asynchronous Method Invocation (AMI) in Python
- The Ice Threading Model

Copyright 2016, ZeroC, Inc.

Example of a File System Server in Python

This page presents the source code for a Python server that implements our file system and communicates with the client we wrote earlier.

The server is remarkably free of code that relates to distribution: most of the server code is simply application logic that would be present just the same as a non-distributed version. Again, this is one of the major advantages of Ice: distribution concerns are kept away from application code so that you can concentrate on developing application logic instead of networking infrastructure.

On this page:

- Implementing a File System Server in Python
- Server Main Program in Python
  - FileI Servant Class in Python
  - DirectoryI Servant Class in Python
  - DirectoryI Data Members in Python
  - DirectoryI Constructor in Python
  - DirectoryI Methods in Python
- Thread Safety in Python

Implementing a File System Server in Python

We have now seen enough of the server-side Python mapping to implement a server for our file system. (You may find it useful to review these Slice definitions before studying the source code.)

Our server is implemented in a single source file, `Server.py`, containing our server's main program as well as the definitions of our `Directory` and `File` servant subclasses.

Server Main Program in Python

Our server main program uses the `Ice.Application` class. The `run` method installs a signal handler, creates an object adapter, instantiates a few servants for the directories and files in the file system, and then activates the adapter. This leads to a main program as follows:

Python

```
import sys, threading, Ice, Filesystem

# DirectoryI servant class ...
# FileI servant class ...

class Server(Ice.Application):
    def run(self, args):
        # Terminate cleanly on receipt of a signal
        #
        self.shutdownOnInterrupt()

        # Create an object adapter (stored in the _adapter
        # static members)
        #
        adapter = self.communicator().createObjectAdapterWithEndpoints(
                        "SimpleFilesystem", "default -p 10000")
        DirectoryI._adapter = adapter
        FileI._adapter = adapter

        # Create the root directory (with name "/" and no parent)
        #
        root = DirectoryI("/", None)

        # Create a file called "README" in the root directory
```


Copyright 2016, ZeroC, Inc.

```
#
file = FileI("README", root)
text = [ "This file system contains a collection of poetry." ]
try:
    file.write(text)
except Filesystem.GenericError, e:
    print e.reason

# Create a directory called "Coleridge"
# in the root directory
#
coleridge = DirectoryI("Coleridge", root)

# Create a file called "Kubla_Khan"
# in the Coleridge directory
#
file = FileI("Kubla_Khan", coleridge)
text = [ "In Xanadu did Kubla Khan",
         "A stately pleasure-dome decree:",
         "Where Alph, the sacred river, ran",
         "Through caverns measureless to man",
         "Down to a sunless sea." ]
try:
    file.write(text)
except Filesystem.GenericError, e:
    print e.reason

# All objects are created, allow client requests now
#
adapter.activate()

# Wait until we are done
#
self.communicator().waitForShutdown()

if self.interrupted():
    print self.appName() + ": terminating"

return 0
```


Copyright 2016, ZeroC, Inc.

```
app = Server()
sys.exit(app.main(sys.argv))
```

The code defines the `Server` class, which derives from `Ice.Application` and contains the main application logic in its `run` method. Much of this code is boiler plate that we saw previously: we create an object adapter, and, towards the end, activate the object adapter and call `waitForShutdown`.

The interesting part of the code follows the adapter creation: here, the server instantiates a few nodes for our file system to create the structure shown below:

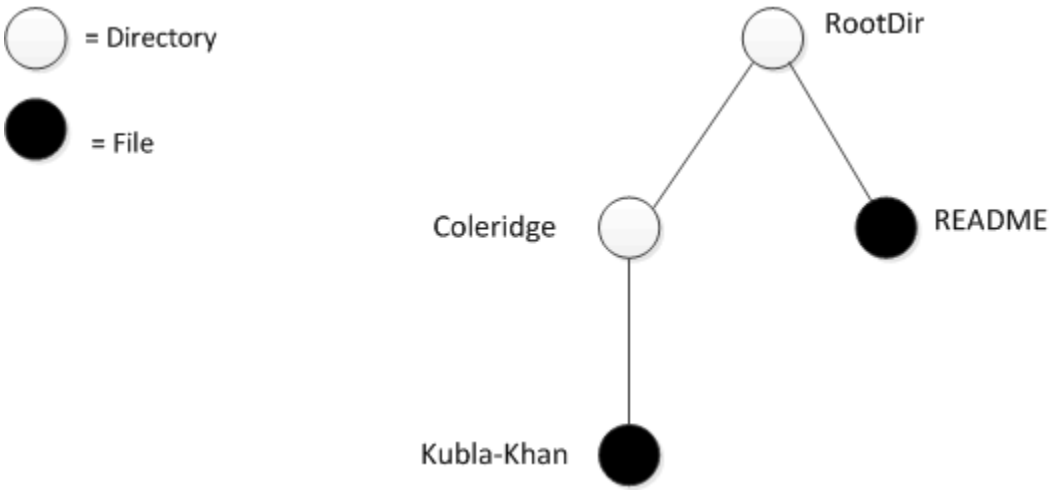

A small file system.

As we will see shortly, the servants for our directories and files are of type `DirectoryI` and `FileI`, respectively. The constructor for either type of servant accepts two parameters, the name of the directory or file to be created and a reference to the servant for the parent directory. (For the root directory, which has no parent, we pass `None`.) Thus, the statement

Python

```
root = DirectoryI("/", None)
```

creates the root directory, with the name `"/"` and no parent directory.

Here is the code that establishes the structure in the above illustration:

Copyright 2016, ZeroC, Inc.

Python

```
# Create the root directory (with name "/" and no parent)
#
root = DirectoryI("/", None)

# Create a file called "README" in the root directory
#
file = FileI("README", root)
text = [ "This file system contains a collection of poetry." ]
try:
    file.write(text)
except Filesystem.GenericError, e:
    print e.reason

# Create a directory called "Coleridge"
# in the root directory
#
coleridge = DirectoryI("Coleridge", root)

# Create a file called "Kubla_Khan"
# in the Coleridge directory
#
file = FileI("Kubla_Khan", coleridge)
text = [ "In Xanadu did Kubla Khan",
         "A stately pleasure-dome decree:",
         "Where Alph, the sacred river, ran",
         "Through caverns measureless to man",
         "Down to a sunless sea." ]
try:
    file.write(text)
except Filesystem.GenericError, e:
    print e.reason
```

We first create the root directory and a file `README` within the root directory. (Note that we pass a reference to the root directory as the parent when we create the new node of type `FileI`.)

The next step is to fill the file with text:

Python

```
text = [ "This file system contains a collection of poetry." ]
try:
    file.write(text)
except Filesystem.GenericError, e:
    print e.reason
```

Recall that Slice sequences map to Python lists. The Slice type `Lines` is simply a list of strings; we add a line of text to our `README` file by initializing the `text` list to contain one element.

Copyright 2016, ZeroC, Inc.

Finally, we call the Slice `write` operation on our `FileI` servant by simply writing:

Python

```
            file.write(text)
```

This statement is interesting: the server code invokes an operation on one of its own servants. Because the call happens via a reference to the servant (of type `FileI`) and not via a proxy (of type `FilePrx`), the Ice run time does not know that this call is even taking place — such a direct call into a servant is not mediated by the Ice run time in any way and is dispatched as an ordinary Python method call.

In similar fashion, the remainder of the code creates a subdirectory called `Coleridge` and, within that directory, a file called `Kubla_Khan` to complete the structure in the above illustration.

`FileI` Servant Class in Python

Our `FileI` servant class has the following basic structure:

Python

```
class FileI(Filesystem.File):
    # Constructor and operations here...

    _adapter = None
```

The class has a number of data members:

- `_adapter`
  This class member stores a reference to the single object adapter we use in our server.
- `_name`
  This instance member stores the name of the file incarnated by the servant.
- `_parent`
  This instance member stores the reference to the servant for the file's parent directory.
- `_lines`
  This instance member holds the contents of the file.

The `_name`, `_parent`, and `_lines` data members are initialized by the constructor:

Copyright 2016, ZeroC, Inc.

Python

```
    def __init__(self, name, parent):
        self._name = name
        self._parent = parent
        self._lines = []

        assert(self._parent != None)

        # Create an identity
        #
        myID = Ice.Identity()
        myID.name = Ice.generateUUID()

        # Add ourself to the object adapter
        #
        self._adapter.add(self, myID)

        # Create a proxy for the new node and
        # add it as a child to the parent
        #
        thisNode = Filesystem.NodePrx.uncheckedCast(self._adapter.createProxy(myID)
)
        self._parent.addChild(thisNode)
```

After initializing the instance members, the code verifies that the reference to the parent is not `None` because every file must have a parent directory. The constructor then generates an identity for the file by calling `Ice.generateUUID` and adds itself to the servant map by calling `ObjectAdapter.add`. Finally, the constructor creates a proxy for this file and calls the `addChild` method on its parent directory. `addChil d` is a helper function that a child directory or file calls to add itself to the list of descendant nodes of its parent directory. We will see the implementation of this function in `DirectoryI` Methods.

The remaining methods of the `FileI` class implement the Slice operations we defined in the `Node` and `File` Slice interfaces:

Python

```
    # Slice Node::name() operation

    def name(self, current=None):
        return self._name

    # Slice File::read() operation

    def read(self, current=None):
        return self._lines

    # Slice File::write() operation

    def write(self, text, current=None):
        self._lines = text
```


Copyright 2016, ZeroC, Inc.

The `name` method is inherited from the generated `Node` class. It simply returns the value of the `_name` instance member.

The `read` and `write` methods are inherited from the generated `File` class and simply return and set the `_lines` instance member.

`DirectoryI` Servant Class in Python

The `DirectoryI` class has the following basic structure:

Python

```
class DirectoryI(Filesystem.Directory):
    # Constructor and operations here...

    _adapter = None
```

`DirectoryI` Data Members in Python

As for the `FileI` class, we have data members to store the object adapter, the name, and the parent directory. (For the root directory, the `_parent` member holds `None`.) In addition, we have a `_contents` data member that stores the list of child directories. These data members are initialized by the constructor:

Python

```
    def __init__(self, name, parent):
        self._name = name
        self._parent = parent
        self._contents = []

        # Create an identity. The
        # parent has the fixed identity "RootDir"
        #
        myID = Ice.Identity()
        if(self._parent):
            myID.name = Ice.generateUUID()
        else:
            myID.name = "RootDir"

        # Add ourself to the object adapter
        #
        self._adapter.add(self, myID)

        # Create a proxy for the new node and
        # add it as a child to the parent
        #
        thisNode = Filesystem.NodePrx.uncheckedCast(self._adapter.createProxy(myID)
)
        if self._parent:
            self._parent.addChild(thisNode)
```

`DirectoryI` Constructor in Python

Copyright 2016, ZeroC, Inc.

The constructor creates an identity for the new directory by calling `Ice.generateUUID`. (For the root directory, we use the fixed identity `"RootDir"`.) The servant adds itself to the servant map by calling `ObjectAdapter.add` and then creates a proxy to itself and passes it to the `addChild` helper function.

`DirectoryI` Methods in Python

`addChild` simply adds the passed reference to the `_contents` list:

Python

```
def addChild(self, child):
    self._contents.append(child)
```

The remainder of the operations, `name` and `list`, are trivial:

Python

```
def name(self, current=None):
    return self._name

def list(self, current=None):
    return self._contents
```

Thread Safety in Python

The server code we have written so far is not quite correct as it stands: if two clients access the same file in parallel, each via a different thread, one thread may read the `_lines` data member while another thread updates it. Obviously, if that happens, we may write or return garbage or, worse, crash the server. However, we can make the `read` and `write` operations thread-safe with a few trivial changes to the `FileI` class:

Copyright 2016, ZeroC, Inc.

Python

```
def __init__(self, name, parent):
    self._name = name
    self._parent = parent
    self._lines = []
    self._mutex = threading.Lock()

    # ...

def name(self, current=None):
    return self._name

def read(self, current=None):
    self._mutex.acquire()
    lines = self._lines[:] # Copy the list
    self._mutex.release()
    return lines

def write(self, text, current=None):
    self._mutex.acquire()
    self._lines = text
    self._mutex.release()
```

We modified the constructor to add the instance member `_mutex`, and then enclosed our `read` and `write` implementations in a critical section. (The `name` method does not require a critical section because the file's name is immutable.)

No changes for thread safety are necessary in the `DirectoryI` class because the `Directory` interface, in its current form, defines no operations that modify the object.

See Also

- Slice for a Simple File System
- Python Mapping for Sequences
- Example of a File System Client in Python
- The Server-Side main Program in Python

Copyright 2016, ZeroC, Inc.

Ruby Mapping

Ice currently provides a client-side mapping for Ruby, but not a server-side mapping.

Topics

- Client-Side Slice-to-Ruby Mapping

Copyright 2016, ZeroC, Inc.

Client-Side Slice-to-Ruby Mapping

The client-side Slice-to-Ruby mapping defines how Slice data types are translated to Ruby types, and how clients invoke operations, pass parameters, and handle errors. Much of the Ruby mapping is intuitive. For example, Slice sequences map to Ruby arrays, so there is essentially nothing new you have to learn in order to use Slice sequences in Ruby.

The Ruby API to the Ice run time is fully thread-safe. Obviously, you must still synchronize access to data from different threads. For example, if you have two threads sharing a sequence, you cannot safely have one thread insert into the sequence while another thread is iterating over the sequence. However, you only need to concern yourself with concurrent access to your own data — the Ice run time itself is fully thread safe, and none of the Ice API calls require you to acquire or release a lock before you safely can make the call.

Much of what appears in this chapter is reference material. We suggest that you skim the material on the initial reading and refer back to specific sections as needed. However, we recommend that you read at least the mappings for exceptions, interfaces, and operations in detail because these sections cover how to call operations from a client, pass parameters, and handle exceptions.

In order to use the Ruby mapping, you should need no more than the Slice definition of your application and knowledge of the Ruby mapping rules. In particular, looking through the generated code in order to discern how to use the Ruby mapping is likely to be inefficient, due to the amount of detail. Of course, occasionally, you may want to refer to the generated code to confirm a detail of the mapping, but we recommend that you otherwise use the material presented here to see how to write your client-side code.

The `Ice` Module

All of the APIs for the Ice run time are nested in the `Ice` module, to avoid clashes with definitions for other libraries or applications. Some of the contents of the `Ice` module are generated from Slice definitions; other parts of the `Ice` module provide special-purpose definitions that do not have a corresponding Slice definition. We will incrementally cover the contents of the `Ice` m odule throughout the remainder of the manual.

A Ruby application can load the Ice run time using the `require` statement:

```
require 'Ice'
```

If the statement executes without error, the Ice run time is loaded and available for use. You can determine the version of the Ice run time you have just loaded by calling the `stringVersion` function:

```
icever = Ice::stringVersion()
```

Topics

- Ruby Mapping for Identifiers
- Ruby Mapping for Modules
- Ruby Mapping for Built-In Types
- Ruby Mapping for Enumerations
- Ruby Mapping for Structures
- Ruby Mapping for Sequences
- Ruby Mapping for Dictionaries
- Ruby Mapping for Constants
- Ruby Mapping for Exceptions
- Ruby Mapping for Interfaces
- Ruby Mapping for Operations
- Ruby Mapping for Classes
- Code Generation in Ruby
- The main Program in Ruby
- Using Slice Checksums in Ruby
- Example of a File System Client in Ruby

Copyright 2016, ZeroC, Inc.

Ruby Mapping for Identifiers

A Slice identifier maps to an identical Ruby identifier. For example, the Slice identifier `Clock` becomes the Ruby identifier `Clock`. There are two exceptions to this rule:

1. Ruby requires the names of classes, modules, and constants to begin with an upper case letter. If a Slice identifier maps to the name of a Ruby class, module, or constant, and the Slice identifier does not begin with an upper case letter, the mapping replaces the leading character with its upper case equivalent. For example, the Slice identifier `bankAccount` is mapped as `BankAccount`.
2. If a Slice identifier is the same as a Ruby keyword, the corresponding Ruby identifier is prefixed with an underscore. For example, the Slice identifier `while` is mapped as `_while`.

You should try to avoid such identifiers as much as possible.

See Also

- Lexical Rules
- Ruby Mapping for Modules
- Ruby Mapping for Built-In Types
- Ruby Mapping for Enumerations
- Ruby Mapping for Structures
- Ruby Mapping for Sequences
- Ruby Mapping for Dictionaries
- Ruby Mapping for Constants
- Ruby Mapping for Exceptions
- Ruby Mapping for Interfaces
- Ruby Mapping for Operations

Copyright 2016, ZeroC, Inc.

Ruby Mapping for Modules

A Slice module maps to a Ruby module with the same name. The mapping preserves the nesting of the Slice definitions.
See Also

- Modules
- Ruby Mapping for Identifiers
- Ruby Mapping for Built-In Types
- Ruby Mapping for Enumerations
- Ruby Mapping for Structures
- Ruby Mapping for Sequences
- Ruby Mapping for Dictionaries
- Ruby Mapping for Constants
- Ruby Mapping for Exceptions
- Ruby Mapping for Interfaces
- Ruby Mapping for Operations

Copyright 2016, ZeroC, Inc.

Ruby Mapping for Built-In Types

On this page:

- Mapping of Slice Built-In Types to Ruby Types
- String Mapping in Ruby

Mapping of Slice Built-In Types to Ruby Types

The Slice built-in types are mapped to Ruby types as shown in this table:

Slice	Ruby
`bool`	`true` or `false`
`byte`	`Fixnum`
`short`	`Fixnum`
`int`	`Fixnum` or `Bignum`
`long`	`Fixnum` or `Bignum`
`float`	`Float`
`double`	`Float`
`string`	`String`

Although Ruby supports arbitrary precision in its integer types, the Ice run time validates integer values to ensure they have valid ranges for their declared Slice types.

String Mapping in Ruby

String values returned as the result of a Slice operation (including return values, out parameters, and data members) contain UTF-8 encoded strings unless the program has installed a string converter, in which case string values use the converter's native encoding instead.

As string input values for a remote Slice operation, Ice accepts `nil` in addition to `String` objects; each occurrence of `nil` is marshaled as an empty string. Ice assumes that all `String` objects contain valid UTF-8 encoded strings unless the program has installed a string converter, in which case Ice assumes that `String` objects use the native encoding expected by the converter.

See Also

- Basic Types
- Ruby Mapping for Identifiers
- Ruby Mapping for Modules
- Ruby Mapping for Enumerations
- Ruby Mapping for Structures
- Ruby Mapping for Sequences
- Ruby Mapping for Dictionaries
- Ruby Mapping for Constants
- Ruby Mapping for Exceptions
- Ruby Mapping for Interfaces
- Ruby Mapping for Operations
- C++ Strings and Character Encoding

Copyright 2016, ZeroC, Inc.

Ruby Mapping for Enumerations

Ruby does not have an enumerated type, so a Slice enumeration is emulated using a Ruby class: the name of the Slice enumeration becomes the name of the Ruby class; for each enumerator, the class contains a constant with the same name as the enumerator. For example:

Slice

```
enum Fruit { Apple, Pear, Orange };
```

The generated Ruby class looks as follows:

Ruby

```
class Fruit
    include Comparable

    Apple = # ...
    Pear = # ...
    Orange = # ...

    def Fruit.from_int(val)

    def to_i

    def to_s

    def <=>(other)

    def hash

    # ...
end
```

The compiler generates a class constant for each enumerator that holds a corresponding instance of `Fruit`. The `from_int` class method returns an instance given its Slice value, while `to_i` returns the Slice value of an enumerator and `to_s` returns its Slice identifier.

Given the above definitions, we can use enumerated values as follows:

Copyright 2016, ZeroC, Inc.

Ruby

```
f1 = Fruit::Apple
f2 = Fruit::Orange

if f1 == Fruit::Apple   # Compare for equality
    # ...

if f1 < f2              # Compare two enums
    # ...

case f2
when Fruit::Orange
    puts "found Orange"
else
    puts "found #{f2.to_s}"
end
```

Comparison operators are available as a result of including `Comparable`, which means a program can compare enumerators according to their Slice values. Note that, when using custom enumerator values, the order of enumerators by their Slice values may not match their order of declaration.

Suppose we modify the Slice definition to include a custom enumerator value:

Slice

```
enum Fruit { Apple, Pear = 3, Orange };
```

We can use `from_int` to examine the Slice values of the enumerators:

Ruby

```
Fruit::from_int(0)  # Apple
Fruit::from_int(1)  # nil
Fruit::from_int(3)  # Pear
Fruit::from_int(4)  # Orange
```

See Also

- Enumerations
- Ruby Mapping for Identifiers
- Ruby Mapping for Modules
- Ruby Mapping for Built-In Types
- Ruby Mapping for Structures
- Ruby Mapping for Sequences
- Ruby Mapping for Dictionaries
- Ruby Mapping for Constants
- Ruby Mapping for Exceptions
- Ruby Mapping for Interfaces
- Ruby Mapping for Operations

Copyright 2016, ZeroC, Inc.

Copyright 2016, ZeroC, Inc.

Ruby Mapping for Structures

A Slice structure maps to a Ruby class with the same name. For each Slice data member, the Ruby class contains a corresponding instance variable as well as accessors to read and write its value. For example, here is our `Employee` structure once more:

Slice

```
struct Employee {
    long number;
    string firstName;
    string lastName;
};
```

The Ruby mapping generates the following definition for this structure:

Ruby

```
class Employee
    def initialize(number=0, firstName='', lastName='')
        @number = number
        @firstName = firstName
        @lastName = lastName
    end

    def hash
        # ...
    end

    def ==
        # ...
    end

    def inspect
        # ...
    end

    attr_accessor :number, :firstName, :lastName
end
```

The constructor initializes each of the instance variables to a default value appropriate for its type:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False

Copyright 2016, ZeroC, Inc.

`sequence`	Null
`dictionary`	Null
`class/interface`	Null

You can also declare different default values for members of primitive and enumerated types.

The compiler generates a definition for the `hash` method, which allows instances to be used as keys in a hash collection. The `hash` method returns a hash value for the structure based on the value of its data members.

The == method returns true if all members of two structures are (recursively) equal.

The `inspect` method returns a string representation of the structure.

See Also

- Structures
- Ruby Mapping for Identifiers
- Ruby Mapping for Modules
- Ruby Mapping for Built-In Types
- Ruby Mapping for Enumerations
- Ruby Mapping for Sequences
- Ruby Mapping for Dictionaries
- Ruby Mapping for Constants
- Ruby Mapping for Exceptions
- Ruby Mapping for Interfaces
- Ruby Mapping for Operations

Copyright 2016, ZeroC, Inc.

Ruby Mapping for Sequences

On this page:

- Mapping Slice Sequences to Ruby Arrays
- Mapping for Byte Sequences in Ruby

Mapping Slice Sequences to Ruby Arrays

A Slice sequence maps to a Ruby array; the only exception is a sequence of bytes, which maps to a string. The use of a Ruby array means that the mapping does not generate a separate named type for a Slice sequence. It also means that you can take advantage of all the array functionality provided by Ruby. For example:

Slice

```
sequence<Fruit> FruitPlatter;
```

We can use the `FruitPlatter` sequence as shown below:

Ruby

```
platter = [ Fruit::Apple, Fruit::Pear ]
platter.push(Fruit::Orange)
```

The Ice run time validates the elements of a sequence to ensure that they are compatible with the declared type; a `TypeError` exception is raised if an incompatible type is encountered.

Mapping for Byte Sequences in Ruby

A Ruby string can contain arbitrary 8-bit binary data, therefore it is a more efficient representation of a byte sequence than a Ruby array in both memory utilization and throughput performance.

When receiving a byte sequence (as the result of an operation, as an out parameter, or as a member of a data structure), the value is always represented as a string. When sending a byte sequence as an operation parameter or data member, the Ice run time accepts both a string and an array of integers as legal values. For example, consider the following Slice definitions:

Slice

```
// Slice
sequence<byte> Data;

interface I {
    void sendData(Data d);
    Data getData();
};
```

The interpreter session below uses these Slice definitions to demonstrate the mapping for a sequence of bytes:

Copyright 2016, ZeroC, Inc.

Ruby

```
> proxy = ...
> proxy.sendData("\0\1\2\3")   # Send as a string
> proxy.sendData([0, 1, 2, 3]) # Send as an array
> d = proxy.getData()
> d.class
=> String
> d
=> "\000\001\002\003"
```

The two invocations of `sendData` are equivalent; however, the second invocation incurs additional overhead as the Ice run time must validate the type and range of each array element.

See Also

- Sequences
- Ruby Mapping for Identifiers
- Ruby Mapping for Modules
- Ruby Mapping for Built-In Types
- Ruby Mapping for Enumerations
- Ruby Mapping for Structures
- Ruby Mapping for Dictionaries
- Ruby Mapping for Constants
- Ruby Mapping for Exceptions
- Ruby Mapping for Interfaces
- Ruby Mapping for Operations

Copyright 2016, ZeroC, Inc.

Ruby Mapping for Dictionaries

Here is the definition of our `EmployeeMap` once more:

Slice

```
dictionary<long, Employee> EmployeeMap;
```

As for sequences, the Ruby mapping does not create a separate named type for this definition. Instead, *all* dictionaries are simply instances of Ruby's hash collection type. For example:

Ruby

```
em = {}

e = Employee.new
e.number = 31
e.firstName = "James"
e.lastName = "Gosling"

em[e.number] = e
```

The Ice run time validates the elements of a dictionary to ensure that they are compatible with the declared type; a `TypeError` exception is raised if an incompatible type is encountered.

See Also

- Dictionaries
- Ruby Mapping for Identifiers
- Ruby Mapping for Modules
- Ruby Mapping for Built-In Types
- Ruby Mapping for Enumerations
- Ruby Mapping for Structures
- Ruby Mapping for Sequences
- Ruby Mapping for Constants
- Ruby Mapping for Exceptions
- Ruby Mapping for Interfaces
- Ruby Mapping for Operations

Copyright 2016, ZeroC, Inc.

Ruby Mapping for Constants

Here are the constant definitions once more:

Slice

```
const bool      AppendByDefault = true;
const byte      LowerNibble = 0x0f;
const string    Advice = "Don't Panic!";
const short     TheAnswer = 42;
const double    PI = 3.1416;

enum Fruit { Apple, Pear, Orange };
const Fruit     FavoriteFruit = Pear;
```

The generated definitions for these constants are shown below:

Ruby

```
AppendByDefault = true
LowerNibble = 15
Advice = "Don't Panic!"
TheAnswer = 42
PI = 3.1416
FavoriteFruit = Fruit::Pear
```

As you can see, each Slice constant is mapped to a Ruby constant with the same name.

Slice string literals that contain non-ASCII characters or universal character names are mapped to Ruby string literals with these characters replaced by their UTF-8 encoding as octal escapes. For example:

Slice

```
const string Egg = "œuf";
const string Heart = "c\u0153ur";
const string Banana = "\U0001F34C";
```

is mapped to:

Ruby

```
Egg = "\305\223uf"
Heart = "c\305\223ur"
Banana = "\360\237\215\214"
```

See Also

- Constants and Literals
- Ruby Mapping for Identifiers

Copyright 2016, ZeroC, Inc.

- Ruby Mapping for Modules
- Ruby Mapping for Built-In Types
- Ruby Mapping for Enumerations
- Ruby Mapping for Structures
- Ruby Mapping for Sequences
- Ruby Mapping for Dictionaries
- Ruby Mapping for Exceptions
- Ruby Mapping for Interfaces
- Ruby Mapping for Operations

Copyright 2016, ZeroC, Inc.

Ruby Mapping for Exceptions

On this page:

- Inheritance Hierarchy for Exceptions in Ruby
- Ruby Mapping for User Exceptions
  - Optional Data Members
- Ruby Mapping for Run-Time Exceptions

Inheritance Hierarchy for Exceptions in Ruby

The mapping for exceptions is based on the inheritance hierarchy shown below:

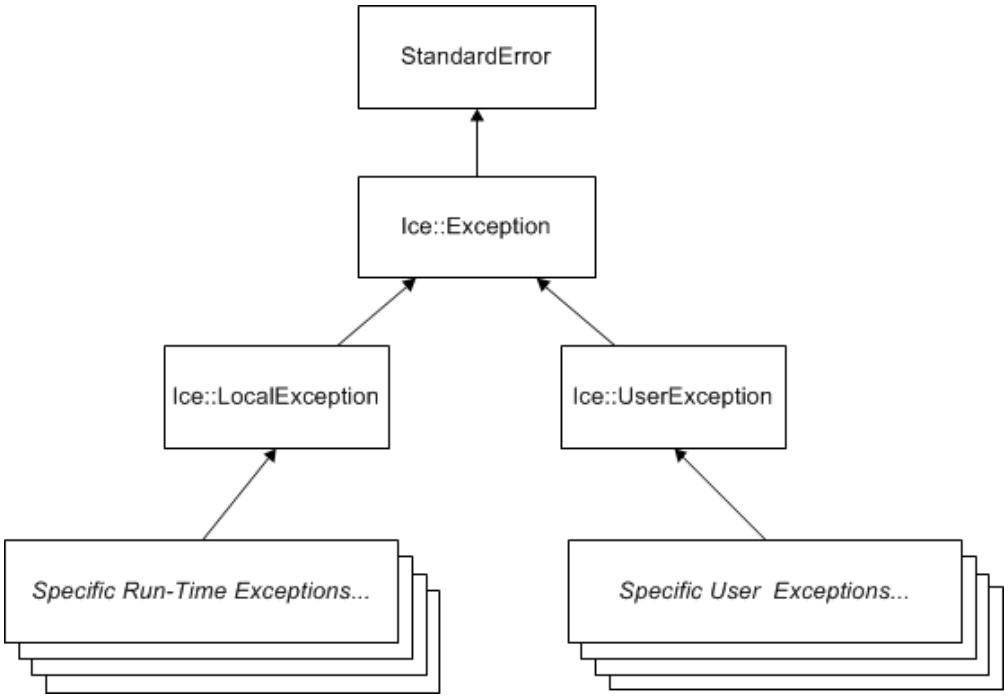

Inheritance structure for Ice exceptions.

The ancestor of all exceptions is `StandardError`, from which `Ice::Exception` is derived. `Ice::LocalException` and `Ice::UserException` are derived from `Ice::Exception` and form the base for all run-time and user exceptions.

Ruby Mapping for User Exceptions

Here is a fragment of the Slice definition for our world time server once more:

Slice

```
exception GenericError {
    string reason;
};
exception BadTimeVal extends GenericError {};
exception BadZoneName extends GenericError {};
```

These exception definitions map to the abbreviated Ruby class definitions shown below:

Copyright 2016, ZeroC, Inc.

Ruby

```
class GenericError < Ice::UserException
    def initialize(reason='')

    def to_s

    def inspect

    attr_accessor :reason
end

class BadTimeVal < GenericError
    def initialize(reason='')

    def to_s

    def inspect
end

class BadZoneName < GenericError
    def initialize(reason='')

    def to_s

    def inspect
end
```

Each Slice exception is mapped to a Ruby class with the same name. The inheritance structure of the Slice exceptions is preserved for the generated classes, so `BadTimeVal` and `BadZoneName` inherit from `GenericError`.

Each exception member corresponds to an instance variable of the instance, which the constructor initializes to a default value appropriate for its type:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	Null
`dictionary`	Null
`class/interface`	Null

You can also declare different default values for members of primitive and enumerated types. For derived exceptions, the constructor has one parameter for each of the base exception's data members, plus one parameter for each of the derived exception's data members, in base-to-derived order. As an example, although `BadTimeVal` and `BadZoneName` do not declare data members, their constructors still accept a value for the inherited data member `reason` in order to pass it to the constructor of the base exception `GenericError`. The generated class defines an accessor for each data member to read and write its value.

Copyright 2016, ZeroC, Inc.

Each exception also defines the standard methods `to_s` and `inspect` to return the name of the exception and a stringified representation of the exception and its members, respectively.

All user exceptions are derived from the base class `Ice::UserException`. This allows you to catch all user exceptions generically by installing a handler for `Ice::UserException`. Similarly, you can catch all Ice run-time exceptions with a handler for `Ice::LocalException`, and you can catch all Ice exceptions with a handler for `Ice::Exception`.

Optional Data Members

Optional data members use the same mapping as required data members, but an optional data member can also be set to the marker value `Ice::Unset` to indicate that the member is unset. A well-behaved program must compare an optional data member to `Ice::Unset` before using the member's value:

Ruby

```
begin
    ...
rescue => ex
    if ex.optionalMember == Ice::Unset
        puts "optionalMember is unset"
    else
        puts "optionalMember = " + ex.optionalMember
    end
end
```

The `Ice::Unset` marker value has different semantics than `nil`. Since `nil` is a legal value for certain Slice types, the Ice run time requires a separate marker value so that it can determine whether an optional value is set. An optional value set to `nil` is considered to be set. If you need to distinguish between an unset value and a value set to `nil`, you can do so as follows:

Ruby

```
begin
    ...
rescue => ex
    if ex.optionalMember == Ice::Unset
        puts "optionalMember is unset"
    elsif ex.optionalMember == nil
        puts "optionalMember is nil"
    else
        puts "optionalMember = " + ex.optionalMember
    end
end
```

Ruby Mapping for Run-Time Exceptions

The Ice run time throws run-time exceptions for a number of pre-defined error conditions. All run-time exceptions directly or indirectly derive from `Ice::LocalException` (which, in turn, derives from `Ice::Exception`).

By catching exceptions at the appropriate point in the inheritance hierarchy, you can handle exceptions according to the category of error they indicate:

- `Ice::LocalException`
  This is the root of the inheritance tree for run-time exceptions.

Copyright 2016, ZeroC, Inc.

- `Ice::UserException`
  This is the root of the inheritance tree for user exceptions.
- `Ice::TimeoutException`
  This is the base exception for both operation-invocation and connection-establishment timeouts.
- `Ice::ConnectTimeoutException`
  This exception is raised when the initial attempt to establish a connection to a server times out.

For example, a `ConnectTimeoutException` can be handled as `ConnectTimeoutException`, `TimeoutException`, `LocalException`, or `Exception`.

You will probably have little need to catch run-time exceptions as their most-derived type and instead catch them as `LocalException`; the fine-grained error handling offered by the remainder of the hierarchy is of interest mainly in the implementation of the Ice run time. Exceptions to this rule are the exceptions related to facet and object life cycles, which you may want to catch explicitly. These exceptions are `FacetNotExistException` and `ObjectNotExistException`, respectively.

See Also

- User Exceptions
- Run-Time Exceptions
- Ruby Mapping for Identifiers
- Ruby Mapping for Modules
- Ruby Mapping for Built-In Types
- Ruby Mapping for Enumerations
- Ruby Mapping for Structures
- Ruby Mapping for Sequences
- Ruby Mapping for Dictionaries
- Ruby Mapping for Constants
- Optional Data Members
- Versioning
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

Ruby Mapping for Interfaces

The mapping of Slice interfaces revolves around the idea that, to invoke a remote operation, you call a member function on a local class instance that is a proxy for the remote object. This makes the mapping easy and intuitive to use because making a remote procedure call is no different from making a local procedure call (apart from error semantics).

On this page:

- Proxy Classes in Ruby
- Ice::ObjectPrx Class in Ruby
- Casting Proxies in Ruby
- Using Proxy Methods in Ruby
- Object Identity and Proxy Comparison in Ruby

Proxy Classes in Ruby

On the client side, a Slice interface maps to a Ruby class with methods that correspond to the operations on those interfaces. Consider the following simple interface:

Slice

```
interface Simple {
    void op();
};
```

The Ruby mapping generates the following definition for use by the client:

Ruby

```
class SimplePrx < Ice::ObjectPrx
    def op(_ctx=nil)
        # ...
    end

    def SimplePrx.ice_staticId()
        # ...
    end
    # ...
end
```

In the client's address space, an instance of `SimplePrx` is the local ambassador for a remote instance of the `Simple` interface in a server and is known as a *proxy instance*. All the details about the server-side object, such as its address, what protocol to use, and its object identity are encapsulated in that instance.

Note that `SimplePrx` inherits from `Ice::ObjectPrx`. This reflects the fact that all Ice interfaces implicitly inherit from `Ice::Object`.

For each operation in the interface, the proxy class has a method of the same name. In the preceding example, we find that the operation `op` has been mapped to the method `op`. Note that `op` accepts an optional trailing parameter `_ctx` representing the operation context. This parameter is a Ruby hash value for use by the Ice run time to store information about how to deliver a request. You normally do not need to use it. (We examine the context parameter in detail in Request Contexts. The parameter is also used by IceStorm.)

Proxy instances are always created on behalf of the client by the Ice run time, so client code never has any need to instantiate a proxy directly.

A value of `nil` denotes the null proxy. The null proxy is a dedicated value that indicates that a proxy points "nowhere" (denotes no object).

Copyright 2016, ZeroC, Inc.

Another method defined by every proxy class is `ice_staticId`, which returns the type ID string corresponding to the interface. As an example, for the Slice interface `Simple` in module `M`, the string returned by `ice_staticId` is `"::M::Simple"`.

`Ice::ObjectPrx` Class in Ruby

All Ice objects have `Object` as the ultimate ancestor type, so all proxies inherit from `Ice::ObjectPrx`. `ObjectPrx` provides a number of methods:

Ruby

```
class ObjectPrx
    def eql?(proxy)
    def ice_getIdentity
    def ice_isA(id)
    def ice_ids
    def ice_id
    def ice_ping
    # ...
end
```

The methods behave as follows:

- **`eql?`**
  The implementation of this standard Ruby method compares two proxies for equality. Note that all aspects of proxies are compared by this operation, such as the communication endpoints for the proxy. This means that, in general, if two proxies compare unequal, that does *not* imply that they denote different objects. For example, if two proxies denote the same Ice object via different transport endpoints, `eql?` returns `false` even though the proxies denote the same object.

- **`ice_getIdentity`**
  This method returns the identity of the object denoted by the proxy. The identity of an Ice object has the following Slice type:

  Slice

  ```
  module Ice {
      struct Identity {
          string name;
          string category;
      };
  };
  ```

  To see whether two proxies denote the same object, first obtain the identity for each object and then compare the identities:

Copyright 2016, ZeroC, Inc.

Ruby

```
proxy1 = ...
proxy2 = ...
id1 = proxy1.ice_getIdentity
id2 = proxy2.ice_getIdentity

if id1 == id2
    # proxy1 and proxy2 denote the same object
else
    # proxy1 and proxy2 denote different objects
end
```

- **ice_isA**
  The ice_isA method determines whether the object denoted by the proxy supports a specific interface. The argument to ice_isA is a type ID. For example, to see whether a proxy of type ObjectPrx denotes a Printer object, we can write:

  Ruby

  ```
  proxy = ...
  if proxy && proxy.ice_isA("::Printer")
      # proxy denotes a Printer object
  else
      # proxy denotes some other type of object
  end
  ```

  Note that we are testing whether the proxy is nil before attempting to invoke the ice_isA method. This avoids getting a run-time error if the proxy is nil.

- **ice_ids**
  The ice_ids method returns an array of strings representing all of the type IDs that the object denoted by the proxy supports.

- **ice_id**
  The ice_id method returns the type ID of the object denoted by the proxy. Note that the type returned is the type of the actual object, which may be more derived than the static type of the proxy. For example, if we have a proxy of type BasePrx, with a static type ID of ::Base, the return value of ice_id might be "::Base", or it might be something more derived, such as "::Derived".

- **ice_ping**
  The ice_ping method provides a basic reachability test for the object. If the object can physically be contacted (that is, the object exists and its server is running and reachable), the call completes normally; otherwise, it throws an exception that indicates why the object could not be reached, such as ObjectNotExistException or ConnectTimeoutException.

The ice_isA, ice_ids, ice_id, and ice_ping methods are remote operations and therefore support an additional overloading that accepts a request context. Also note that there are other methods in ObjectPrx, not shown here. These methods provide different ways to dispatch a call and also provide access to an object's facets.

Casting Proxies in Ruby

The Ruby mapping for a proxy also generates two class methods:

Copyright 2016, ZeroC, Inc.

Ruby

```
class SimplePrx < Ice::ObjectPrx
    # ...

    def SimplePrx.checkedCast(proxy, facet='', ctx={})

    def SimplePrx.uncheckedCast(proxy, facet='')
end
```

The method names `checkedCast` and `uncheckedCast` are reserved for use in proxies. If a Slice interface defines an operation with either of those names, the mapping escapes the name in the generated proxy by prepending an underscore. For example, an interface that defines an operation named `checkedCast` is mapped to a proxy with a method named `_checkedCast`.

For `checkedCast`, if the passed proxy is for an object of type `Simple`, or a proxy for an object with a type derived from `Simple`, the cast returns a reference to a proxy of type `SimplePrx`; otherwise, if the passed proxy denotes an object of a different type (or if the passed proxy is `nil`), the cast returns `nil`.

Given a proxy of any type, you can use a checkedCast to determine whether the corresponding object supports a given type, for example:

Ruby

```
obj = ...       # Get a proxy from somewhere...

simple = SimplePrx::checkedCast(obj)
if simple
    # Object supports the Simple interface...
else
    # Object is not of type Simple...
end
```

Note that a `checkedCast` contacts the server. This is necessary because only the server implementation has definite knowledge of the type of an object. As a result, a `checkedCast` may throw a `ConnectTimeoutException` or an `ObjectNotExistException`.

In contrast, an `uncheckedCast` does not contact the server and unconditionally returns a proxy of the requested type. However, if you do use an `uncheckedCast`, you must be certain that the proxy really does support the type you are casting to; otherwise, if you get it wrong, you will most likely get a run-time exception when you invoke an operation on the proxy. The most likely error for such a type mismatch is `OperationNotExistException`. However, other exceptions, such as a marshaling exception are possible as well. And, if the object happens to have an operation with the correct name, but different parameter types, no exception may be reported at all and you simply end up sending the invocation to an object of the wrong type; that object may do rather nonsensical things. To illustrate this, consider the following two interfaces:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Process {
    void launch(int stackSize, int dataSize);
};

// ...

interface Rocket {
    void launch(float xCoord, float yCoord);
};
```

Suppose you expect to receive a proxy for a `Process` object and use an `uncheckedCast` to down-cast the proxy:

Ruby

```
obj = ...                                # Get proxy...
process = ProcessPrx::uncheckedCast(obj) # No worries...
process.launch(40, 60)                   # Oops...
```

If the proxy you received actually denotes a `Rocket` object, the error will go undetected by the Ice run time: because `int` and `float` have the same size and because the Ice protocol does not tag data with its type on the wire, the implementation of `Rocket::launch` will simply misinterpret the passed integers as floating-point numbers.

In fairness, this example is somewhat contrived. For such a mistake to go unnoticed at run time, both objects must have an operation with the same name and, in addition, the run-time arguments passed to the operation must have a total marshaled size that matches the number of bytes that are expected by the unmarshaling code on the server side. In practice, this is extremely rare and an incorrect `uncheckedCast` typically results in a run-time exception.

Using Proxy Methods in Ruby

The base proxy class `ObjectPrx` supports a variety of methods for customizing a proxy. Since proxies are immutable, each of these "factory methods" returns a copy of the original proxy that contains the desired modification. For example, you can obtain a proxy configured with a ten second timeout as shown below:

Ruby

```
proxy = communicator.stringToProxy(...)
proxy = proxy.ice_timeout(10000)
```

A factory method returns a new proxy object if the requested modification differs from the current proxy, otherwise it returns the current proxy. With few exceptions, factory methods return a proxy of the same type as the current proxy, therefore it is generally not necessary to repeat a down-cast after using a factory method. The example below demonstrates these semantics:

Copyright 2016, ZeroC, Inc.

Ruby

```
base = communicator.stringToProxy(...)
hello = Demo::HelloPrx::checkedCast(base)
hello = hello.ice_timeout(10000) # Type is not discarded
hello.sayHello()
```

The only exceptions are the factory methods `ice_facet` and `ice_identity`. Calls to either of these methods may produce a proxy for an object of an unrelated type, therefore they return a base proxy that you must subsequently down-cast to an appropriate type.

Object Identity and Proxy Comparison in Ruby

Proxy objects support comparison using the comparison operators ==, `!=`, and <=>, as well as the `eql?` method. Note that proxy comparison uses *all* of the information in a proxy for the comparison. This means that not only the object identity must match for a comparison to succeed, but other details inside the proxy, such as the protocol and endpoint information, must be the same. In other words, comparison tests for *proxy* identity, *not* object identity. A common mistake is to write code along the following lines:

Ruby

```
p1 = ...        # Get a proxy...
p2 = ...        # Get another proxy...

if p1 != p2
    # p1 and p2 denote different objects       # WRONG!
else
    # p1 and p2 denote the same object         # Correct
end
```

Even though `p1` and `p2` differ, they may denote the same Ice object. This can happen because, for example, both `p1` and `p2` embed the same object identity, but each uses a different protocol to contact the target object. Similarly, the protocols may be the same, but denote different endpoints (because a single Ice object can be contacted via several different transport endpoints). In other words, if two proxies compare equal, we know that the two proxies denote the same object (because they are identical in all respects); however, if two proxies compare unequal, we know absolutely nothing: the proxies may or may not denote the same object.

To compare the object identities of two proxies, you can use helper functions in the `Ice` module:

Ruby

```
def proxyIdentityCompare(lhs, rhs)
def proxyIdentityAndFacetCompare(lhs, rhs)
```

`proxyIdentityCompare` allows you to correctly compare proxies for identity:

Copyright 2016, ZeroC, Inc.

Ruby

```
p1 = ...        # Get a proxy...
p2 = ...        # Get another proxy...

if Ice.proxyIdentityCompare(p1, p2) != 0
    # p1 and p2 denote different objects       # Correct
else
    # p1 and p2 denote the same object         # Correct
end
```

The function returns `0` if the identities are equal, `-1` if `p1` is less than `p2`, and `1` if `p1` is greater than `p2`. (The comparison uses `name` as the major sort key and `category` as the minor sort key.)

The `proxyIdentityAndFacetCompare` function behaves similarly, but compares both the identity and the facet name.

See Also

- Interfaces, Operations, and Exceptions
- Proxies for Ice Objects
- Type IDs
- Ruby Mapping for Operations
- Request Contexts
- Operations on Object
- Proxy Methods
- Versioning
- IceStorm

Copyright 2016, ZeroC, Inc.

Ruby Mapping for Operations

On this page:

- Basic Ruby Mapping for Operations
- Normal and idempotent Operations in Ruby
- Passing Parameters in Ruby
  - In-Parameters in Ruby
  - Out-Parameters in Ruby
  - Parameter Type Mismatches in Ruby
  - Null Parameters in Ruby
  - Optional Parameters in Ruby
- Exception Handling in Ruby

Basic Ruby Mapping for Operations

As we saw in the Ruby mapping for interfaces, for each operation on an interface, the proxy class contains a corresponding method with the same name. To invoke an operation, you call it via the proxy. For example, here is part of the definitions for our file system:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };
    // ...
};
```

The `name` operation returns a value of type `string`. Given a proxy to an object of type `Node`, the client can invoke the operation as follows:

Ruby

```
node = ...          # Initialize proxy
name = node.name()  # Get name via RPC
```

Normal and `idempotent` Operations in Ruby

You can add an `idempotent` qualifier to a Slice operation. As far as the signature for the corresponding proxy method is concerned, idempotent has no effect. For example, consider the following interface:

Slice

```
interface Example {
                string op1();
    idempotent  string op2();
};
```

The proxy class for this is:

Copyright 2016, ZeroC, Inc.

Ruby

```
class ExamplePrx < Ice::ObjectPrx
    def op1(_ctx=nil)

    def op2(_ctx=nil)
end
```

Because `idempotent` affects an aspect of call dispatch, not interface, it makes sense for the two methods to look the same.

Passing Parameters in Ruby

In-Parameters in Ruby

All parameters are passed by reference in the Ruby mapping; it is guaranteed that the value of a parameter will not be changed by the invocation.

Here is an interface with operations that pass parameters of various types from client to server:

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ClientToServer {
    void op1(int i, float f, bool b, string s);
    void op2(NumberAndString ns, StringSeq ss, StringTable st);
    void op3(ClientToServer* proxy);
};
```

The Slice compiler generates the following proxy for this definition:

Copyright 2016, ZeroC, Inc.

Ruby

```
class ClientToServerPrx < Ice::ObjectPrx
    def op1(i, f, b, s, _ctx=nil)

    def op2(ns, ss, st, _ctx=nil)

    def op3(proxy, _ctx=nil)
end
```

Given a proxy to a `ClientToServer` interface, the client code can pass parameters as in the following example:

Ruby

```
p = ...                                 # Get proxy...

p.op1(42, 3.14, true, "Hello world!")   # Pass simple literals

i = 42
f = 3.14
b = true
s = "Hello world!"
p.op1(i, f, b, s)                       # Pass simple variables

ns = NumberAndString.new()
ns.x = 42
ns.str = "The Answer"
ss = [ "Hello world!" ]
st = {}
st[0] = ns
p.op2(ns, ss, st)                       # Pass complex variables

p.op3(p)                                # Pass proxy
```

Out-Parameters in Ruby

As in Java, Ruby functions do not support reference arguments. That is, it is not possible to pass an uninitialized variable to a Ruby function in order to have its value initialized by the function. The Java mapping overcomes this limitation with the use of *holder classes* that represent each `out` parameter. The Ruby mapping takes a different approach, one that is more natural for Ruby users.

The semantics of `out` parameters in the Ruby mapping depend on whether the operation returns one value or multiple values. An operation returns multiple values when it has declared multiple `out` parameters, or when it has declared a non-`void` return type and at least one `out` parameter.

If an operation returns multiple values, the client receives them in the form of a *result array*. A non-`void` return value, if any, is always the first element in the result array, followed by the `out` parameters in the order of declaration.

If an operation returns only one value, the client receives the value itself.

Here again are the same Slice definitions we saw earlier, but this time with all parameters being passed in the `out` direction:

Copyright 2016, ZeroC, Inc.

Slice

```
struct NumberAndString {
    int x;
    string str;
};

sequence<string> StringSeq;

dictionary<long, StringSeq> StringTable;

interface ServerToClient {
    int op1(out float f, out bool b, out string s);
    void op2(out NumberAndString ns,
             out StringSeq ss,
             out StringTable st);
    void op3(out ServerToClient* proxy);
};
```

The Ruby mapping generates the following code for this definition:

Ruby

```
class ClientToServerPrx < Ice::ObjectPrx
    def op1(_ctx=nil)

    def op2(_ctx=nil)

    def op3(_ctx=nil)
end
```

Given a proxy to a `ServerToClient` interface, the client code can receive the results as in the following example:

Ruby

```
p = ...              # Get proxy...
i, f, b, s = p.op1()
ns, ss, st = p.op2()
stcp = p.op3()
```

The operations have no `in` parameters, therefore no arguments are passed to the proxy methods. Since `op1` and `op2` return multiple values, their result arrays are unpacked into separate values, whereas the return value of `op3` requires no unpacking.

Parameter Type Mismatches in Ruby

Although the Ruby compiler cannot check the types of arguments passed to a method, the Ice run time does perform validation on the arguments to a proxy invocation and reports any type mismatches as a `TypeError` exception.

Copyright 2016, ZeroC, Inc.

Null Parameters in Ruby

Some Slice types naturally have "empty" or "not there" semantics. Specifically, sequences, dictionaries, and strings all can be `nil`, but the corresponding Slice types do not have the concept of a null value. To make life with these types easier, whenever you pass `nil` as a parameter or return value of type sequence, dictionary, or string, the Ice run time automatically sends an empty sequence, dictionary, or string to the receiver.

This behavior is useful as a convenience feature: especially for deeply-nested data types, members that are sequences, dictionaries, or strings automatically arrive as an empty value at the receiving end. This saves you having to explicitly initialize, for example, every string element in a large sequence before sending the sequence in order to avoid a run-time error. Note that using null parameters in this way does *not* create null semantics for Slice sequences, dictionaries, or strings. As far as the object model is concerned, these do not exist (only *empty* sequences, dictionaries, and strings do). For example, it makes no difference to the receiver whether you send a string as `nil` or as an empty string: either way, the receiver sees an empty string.

Optional Parameters in Ruby

Optional parameters use the same mapping as required parameters. The only difference is that `Ice::Unset` can be passed as the value of an optional parameter or return value. Consider the following operation:

Slice

```
optional(1) int execute(optional(2) string params, out optional(3) float value);
```

A client can invoke this operation as shown below:

Ruby

```
i, v = proxy.execute("--file log.txt")
i, v = proxy.execute(Ice::Unset)

if v != Ice::Unset
    puts "value = " + v.to_s
end
```

A well-behaved program must always compare an optional parameter to `Ice::Unset` prior to using its value. Keep in mind that the `Ice::Unset` marker value has different semantics than `nil`. Since `nil` is a legal value for certain Slice types, the Ice run time requires a separate marker value so that it can determine whether an optional parameter is set. An optional parameter set to `nil` is considered to be set. If you need to distinguish between an unset parameter and a parameter set to `nil`, you can do so as follows:

Ruby

```
if optionalParam == Ice::Unset
    puts "optionalParam is unset"
elsif optionalParam == nil
    puts "optionalParam is None"
else
    puts "optionalParam = " + optionalParam
end
```


Copyright 2016, ZeroC, Inc.

Exception Handling in Ruby

Any operation invocation may throw a run-time exception and, if the operation has an exception specification, may also throw user exceptions. Suppose we have the following simple interface:

Slice

```
exception Tantrum {
    string reason;
};

interface Child {
    void askToCleanUp() throws Tantrum;
};
```

Slice exceptions are thrown as Ruby exceptions, so you can simply enclose one or more operation invocations in a `begin-rescue` block:

Ruby

```
child = ...        # Get child proxy...

begin
    child.askToCleanUp()
rescue Tantrum => t
    puts "The child says: #{t.reason}"
end
```

Typically, you will catch only a few exceptions of specific interest around an operation invocation; other exceptions, such as unexpected run-time errors, will usually be handled by exception handlers higher in the hierarchy. For example:

Copyright 2016, ZeroC, Inc.

Ruby

```
def run()
    child = ...        # Get child proxy...
    begin
        child.askToCleanUp()
    rescue Tantrum => t
        puts "The child says: #{t.reason}"
        child.scold()  # Recover from error...
    end
    child.praise()     # Give positive feedback...
end

begin
    # ...
    run()
    # ...
rescue Ice::Exception => ex
    print ex.backtrace.join("\n")
end
```

This code handles a specific exception of local interest at the point of call and deals with other exceptions generically. (This is also the strategy we used for our first simple application in Hello World Application.)

See Also

- Operations
- Hello World Application
- Slice for a Simple File System
- Ruby Mapping for Operations
- Ruby Mapping for Interfaces
- Ruby Mapping for Exceptions

Copyright 2016, ZeroC, Inc.

Ruby Mapping for Classes

On this page:

- Basic Ruby Mapping for Classes
- Inheritance from Ice::Object in Ruby
- Class Data Members in Ruby
- Class Constructors in Ruby
- Class Operations in Ruby
- Receiving Objects in Ruby
- Class Factories in Ruby

Basic Ruby Mapping for Classes

A Slice class maps to a Ruby class with the same name. For each Slice data member, the generated class contains an instance variable and accessors to read and write it, just as for structures and exceptions. Consider the following class definition:

Slice

```
class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Ruby mapping generates the following code for this definition:

Copyright 2016, ZeroC, Inc.

Ruby

```
module TimeOfDay_mixin
    include ::Ice::Object_mixin

    # ...

    def inspect
        # ...
    end

    #
    # Operation signatures.
    #
    # def format()

    attr_accessor :hours, :minutes, :seconds
end
class TimeOfDay
    include TimeOfDay_mixin

    def initialize(hour=0, minute=0, second=0)
        @hour = hour
        @minute = minute
        @second = second
    end

    def TimeOfDay.ice_staticId()
        '::M::TimeOfDay'
    end

    # ...
end
```

There are a number of things to note about the generated code:

1. The generated class `TimeOfDay` includes the mixin module `TimeOfDay_mixin`, which in turn includes `Ice::Object_mixin`. This reflects the semantics of Slice classes in that all classes implicitly inherit from `Object`, which is the ultimate ancestor of all classes. Note that `Object` is *not* the same as `Ice::ObjectPrx`. In other words, you *cannot* pass a class where a proxy is expected and vice versa.
2. The constructor defines an instance variable for each Slice data member.
3. The class defines the class method `ice_staticId`.
4. A comment summarizes the method signatures for each Slice operation.

There is quite a bit to discuss here, so we will look at each item in turn.

Inheritance from `Ice::Object` in Ruby

In other language mappings, the inheritance relationship between `Object` and a user-defined Slice class is stated explicitly, in that the generated class derives from a language-specific representation of `Object`. Although its class type allows single inheritance, Ruby's loosely-typed nature places less emphasis on class hierarchies and relies more on *duck typing*.

In Ruby, an object's type is typically less important than the methods it supports. *If it looks like a duck, and acts like a duck, then it*

Copyright 2016, ZeroC, Inc.

is a duck.

The Slice mapping for a class follows this convention by placing most of the necessary machinery in a mixin module that the generated class includes into its definition. The Ice run time requires an instance of a Slice class to include the mixin module and define values for the declared data members, but does not require that the object be an instance of the generated class.

As shown in the illustration below, classes have no relationship to `Ice::ObjectPrx` (which is at the base of the inheritance hierarchy for proxies), therefore you cannot pass a class where a proxy is expected (and vice versa).

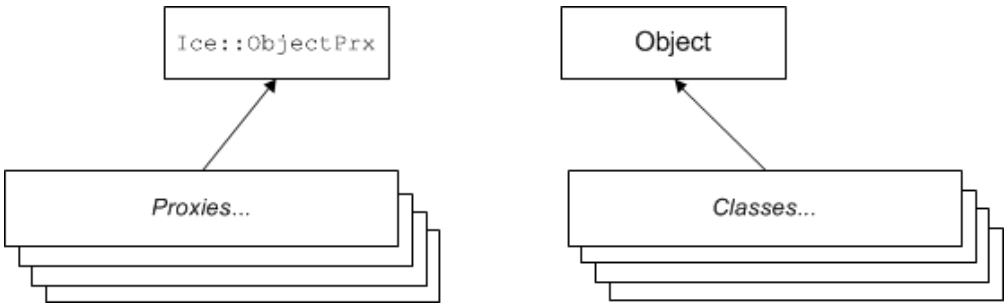

Inheritance from Ice::ObjectPrx and Object.

An instance of a Slice class `C` supports a number of methods:

Ruby

```
def ice_isA(id, current=nil)

def ice_ping(current=nil)

def ice_ids(current=nil)

def ice_id(current=nil)

def C.ice_staticId()

def ice_preMarshal()

def ice_postUnmarshal()
```

The methods behave as follows:

- `ice_isA`
  This method returns `true` if the object supports the given type ID, and `false` otherwise.
- `ice_ping`
  As for interfaces, `ice_ping` provides a basic reachability test for the object.
- `ice_ids`
  This method returns a string sequence representing all of the type IDs supported by this object, including `::Ice::Object`.
- `ice_id`
  This method returns the actual run-time type ID of the object. If you call `ice_id` through a reference to a base instance, the returned type id is the actual (possibly more derived) type ID of the instance.
- `ice_staticId`
  This method returns the static type ID of the class.
- `ice_preMarshal`
  If the object supports this method, the Ice run time invokes it just prior to marshaling the object's state, providing the opportunity for

Copyright 2016, ZeroC, Inc.

the object to validate its declared data members.

- `ice_postUnmarshal`
  If the object supports this method, the Ice run time invokes it after unmarshaling the object's state. An object typically defines this method when it needs to perform additional initialization using the values of its declared data members.

The mixin module `Ice::Object_mixin` supplies default definitions of `ice_isA` and `ice_ping`. For each Slice class, the generated mixin module defines `ice_ids` and `ice_id`, and the generated class defines the `ice_staticId` method.

Note that neither `Ice::Object` nor the generated class override `hash` and `==`, so the default implementations apply.

Class Data Members in Ruby

By default, data members of classes are mapped exactly as for structures and exceptions: for each data member in the Slice definition, the generated class contains a corresponding instance variable and accessor methods.

Optional data members use the same mapping as required data members, but an optional data member can also be set to the marker value `Ice::Unset` to indicate that the member is unset. A well-behaved program must compare an optional data member to `Ice::Unset` before using the member's value:

Ruby

```
v = ...
if v.optionalMember == Ice::Unset
    puts "optionalMember is unset"
else
    puts "optionalMember = " + v.optionalMember
end
```

The `Ice::Unset` marker value has different semantics than `nil`. Since `nil` is a legal value for certain Slice types, the Ice run time requires a separate marker value so that it can determine whether an optional value is set. An optional value set to `nil` is considered to be set. If you need to distinguish between an unset value and a value set to `nil`, you can do so as follows:

Ruby

```
v = ...
if v.optionalMember == Ice::Unset
    puts "optionalMember is unset"
elsif v.optionalMember == nil
    puts "optionalMember is nil"
else
    puts "optionalMember = " + v.optionalMember
end
```

If you wish to restrict access to a data member, you can modify its visibility using the `protected` metadata directive. The presence of this directive causes the Slice compiler to generate the data member with protected visibility. As a result, the member can be accessed only by the class itself or by one of its subclasses. For example, the `TimeOfDay` class shown below has the `protected` metadata directive applied to each of its data members:

Copyright 2016, ZeroC, Inc.

Slice

```
class TimeOfDay {
    ["protected"] short hour;   // 0 - 23
    ["protected"] short minute; // 0 - 59
    ["protected"] short second; // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

The Slice compiler produces the following generated code for this definition:

Ruby

```
module TimeOfDay_mixin
    include ::Ice::Object_mixin

    # ...

    #
    # Operation signatures.
    #
    # def format()

    attr_accessor :hours, :minutes, :seconds
    protected :hours, :hours=
    protected :minutes, :minutes=
    protected :seconds, :seconds=
end
class TimeOfDay
    include TimeOfDay_mixin

    def initialize(hour=0, minute=0, second=0)
        @hour = hour
        @minute = minute
        @second = second
    end

    # ...
end
```

For a class in which all of the data members are protected, the metadata directive can be applied to the class itself rather than to each member individually. For example, we can rewrite the `TimeOfDay` class as follows:

Copyright 2016, ZeroC, Inc.

Slice

```
["protected"] class TimeOfDay {
    short hour;         // 0 - 23
    short minute;       // 0 - 59
    short second;       // 0 - 59
    string format();    // Return time as hh:mm:ss
};
```

Class Constructors in Ruby

Classes have a constructor that assigns to each data member a default value appropriate for its type:

Data Member Type	Default Value
`string`	Empty string
`enum`	First enumerator in enumeration
`struct`	Default-constructed value
Numeric	Zero
`bool`	False
`sequence`	Null
`dictionary`	Null
`class/interface`	Null

You can also declare different default values for data members of primitive and enumerated types.

For derived classes, the constructor has one parameter for each of the base class's data members, plus one parameter for each of the derived class's data members, in base-to-derived order.

Pass the marker value `Ice::Unset` as the value of any optional data members that you wish to be unset.

Class Operations in Ruby

Operations of classes are mapped to methods in the generated class. This means that, if a class contains operations (such as the `format` operation of our `TimeOfDay` class), objects representing instances of `TimeOfDay` must define equivalent methods. For example:

Ruby

```
class TimeOfDayI < TimeOfDay
    def format(current=nil)
        sprintf("%02d:%02d:%02d", @hour, @minute, @second)
    end
end
```

In this case our implementation class `TimeOfDayI` derives from the generated class `TimeOfDay`. An alternative is to include the generated mixin module, which makes it possible for the class to derive from a different base class if necessary:

Copyright 2016, ZeroC, Inc.

Ruby

```
class TimeOfDayI < SomeOtherClass
    include TimeOfDay_mixin

    def format(current=nil)
        sprintf("%02d:%02d:%02d", @hour, @minute, @second)
    end
end
```

As explained earlier, an implementation of a Slice class must include the mixin module but is not required to derive from the generated class.

Ruby allows an existing class to be reopened in order to augment or replace its functionality. This feature provides another way for us to implement a Slice class: reopen the generated class and define the necessary methods:

Ruby

```
class TimeOfDay
    def format(current=nil)
        sprintf("%02d:%02d:%02d", @hour, @minute, @second)
    end
end
```

As an added benefit, this strategy eliminates the need to define a class factory. The next section describes this subject in more detail.

A Slice class such as `TimeOfDay` that declares or inherits an operation is inherently abstract. Ruby does not support the notion of abstract classes or abstract methods, therefore the mapping merely summarizes the required method signatures in a comment for your convenience.

You may notice that the mapping for an operation adds an optional trailing parameter named `current`. For now, you can ignore this parameter and pretend it does not exist.

Receiving Objects in Ruby

We have discussed the ways you can implement a Slice class, but we also need to examine the semantics of receiving an object as the return value or as an out-parameter from an operation invocation. Consider the following simple interface:

Slice

```
interface Time {
    TimeOfDay get();
};
```

When a client invokes the `get` operation, the Ice run time must instantiate and return an instance of the `TimeOfDay` class. Unless we tell it otherwise, the Ice run time in Ruby does exactly that: it instantiates the generated class `TimeOfDay`. Although `TimeOfDay` is logically an abstract class because its Slice equivalent defined an operation, Ruby has no notion of abstract classes and therefore it is legal to create an instance of this class. Furthermore, there are situations in which this is exactly the behavior you want:

- when you have reopened the generated class to define its operations, or
- when your program uses only the data members of an object and does not invoke any of its operations.

On the other hand, if you have defined a Ruby class that implements the Slice class, you need the Ice run time to return an instance of your

Copyright 2016, ZeroC, Inc.

class and not an instance of the generated class. The Ice run time cannot magically know about your implementation class, therefore you must inform the Ice run time by installing a class factory.

Class Factories in Ruby

The Ice run time invokes a class factory when it needs to instantiate an object of a particular type. If no factory is found, the Ice run time instantiates the generated class as described above. To install a factory, we use operations provided by the `Ice::Communicator` interface :

Slice

```
module Ice {
    local interface ObjectFactory {
        Object create(string type);
        void destroy();
    };

    local interface Communicator {
        void addObjectFactory(ObjectFactory factory,

string id);
        ObjectFactory findObjectFactory(string id);
        // ...
    };
};
```

To supply the Ice run time with a factory for our `TimeOfDayI` class, we must create an object that supports the `Ice::ObjectFactory` interface:

Ruby

```
class ObjectFactory
    def create(type)
        fail unless type == M::TimeOfDay::ice_staticId()
        TimeOfDayI.new
    end

    def destroy
        # Nothing to do
    end
end
```

The object factory's `create` method is called by the Ice run time when it needs to instantiate a `TimeOfDay` class. The factory's `destroy` method is called by the Ice run time when its communicator is destroyed.

The `create` method is passed the type ID of the class to instantiate. For our `TimeOfDay` class, the type ID is `"::M::TimeOfDay"`. Our implementation of `create` checks the type ID: if it matches, the method instantiates and returns a `TimeOfDayI` object. For other type IDs, the method fails because it does not know how to instantiate other types of objects.

Note that we used the `ice_staticId` method to obtain the type ID rather than embedding a literal string. Using a literal type ID string in your code is discouraged because it can lead to errors that are only detected at run time. For example, if a Slice class or one of its enclosing modules is renamed and the literal string is not changed accordingly, a receiver will fail to unmarshal the object and the Ice run time will raise `NoObjectFactoryException`. By using `ice_staticId` instead, we avoid any risk of a misspelled or obsolete type ID, and we can discover at compile time if a Slice class or module has been renamed.

Copyright 2016, ZeroC, Inc.

Given a factory implementation, such as our `ObjectFactory`, we must inform the Ice run time of the existence of the factory:

Ruby

```
ic = ...   # Get Communicator...
ic.addObjectFactory(ObjectFactory.new, M::TimeOfDay::ice_staticId())
```

Now, whenever the Ice run time needs to instantiate a class with the type ID `"::M::TimeOfDay"`, it calls the `create` method of the registered `ObjectFactory` instance.

The `destroy` operation of the object factory is invoked by the Ice run time when the communicator is destroyed. This gives you a chance to clean up any resources that may be used by your factory. Do not call `destroy` on the factory while it is registered with the communicator — if you do, the Ice run time has no idea that this has happened and, depending on what your `destroy` implementation is doing, may cause undefined behavior when the Ice run time tries to next use the factory.

The run time guarantees that `destroy` will be the last call made on the factory, that is, `create` will not be called concurrently with `destroy`, and `create` will not be called once `destroy` has been called. However, calls to `create` can be made concurrently.

Note that you cannot register a factory for the same type ID twice: if you call `addObjectFactory` with a type ID for which a factory is registered, the Ice run time throws an `AlreadyRegisteredException`.

Finally, keep in mind that if a class has only data members, but no operations, you need not create and register an object factory to transmit instances of such a class.

See Also

- Classes
- Type IDs
- Optional Data Members
- The Current Object

Copyright 2016, ZeroC, Inc.

Code Generation in Ruby

The Ruby mapping supports two forms of code generation: dynamic and static.

On this page:

- Dynamic Code Generation in Ruby
  - Ice::loadSlice Options in Ruby
  - Locating Slice Files in Ruby
  - Loading Multiple Slice Files in Ruby
  - Limitations of Dynamic Code Generation in Ruby
- Static Code Generation in Ruby
  - Compiler Output in Ruby
  - Include Files in Ruby
- Static Versus Dynamic Code Generation in Ruby
  - Application Considerations for Code Generation in Ruby
  - Mixing Static and Dynamic Generation in Ruby
- slice2rb Command-Line Options

Dynamic Code Generation in Ruby

Using dynamic code generation, Slice files are "loaded" at run time and dynamically translated into Ruby code, which is immediately compiled and available for use by the application. This is accomplished using the `Ice::loadSlice` method, as shown in the following example:

Ruby

```
Ice::loadSlice("Color.ice")
puts "My favorite color is #{M::Color.blue.to_s}"
```

For this example, we assume that `Color.ice` contains the following definitions:

Slice

```
module M {
    enum Color { red, green, blue };
};
```

`Ice::loadSlice` Options in Ruby

The `Ice::loadSlice` method behaves like a Slice compiler in that it accepts command-line arguments for specifying preprocessor options and controlling code generation. The arguments must include at least one Slice file.

The function has the following Ruby definition:

Ruby

```
def loadSlice(cmd, args=[])
```

The command-line arguments can be specified entirely in the first argument, `cmd`, which must be a string. The optional second argument can be used to pass additional command-line arguments as a list; this is useful when the caller already has the arguments in list form. The function always returns `nil`.

Copyright 2016, ZeroC, Inc.

For example, the following calls to `Ice::loadSlice` are functionally equivalent:

Ruby

```
Ice::loadSlice("-I/opt/IceRuby/slice Color.ice")
Ice::loadSlice("-I/opt/IceRuby/slice", ["Color.ice"])
Ice::loadSlice("", ["-I/opt/IceRuby/slice", "Color.ice"])
```

In addition to the standard compiler options, `Ice::loadSlice` also supports the following command-line options:

- `--all`
  Generate code for all Slice definitions, including those from included files.
- `--checksum`
  Generate checksums for Slice definitions.

Locating Slice Files in Ruby

If your Slice files depend on Ice types, you can avoid hard-coding the path name of your Ice installation directory by calling the `Ice::getSliceDir` function:

Ruby

```
Ice::loadSlice("-I" + Ice::getSliceDir() + " Color.ice")
```

This function attempts to locate the `slice` subdirectory of your Ice installation using an algorithm that succeeds for the following scenarios:

- Installation of a binary Ice archive
- Installation of an Ice source distribution using `make install`
- Installation via a Windows installer
- RPM installation on Linux
- Execution inside a compiled Ice source distribution

If the `slice` subdirectory can be found, `getSliceDir` returns its absolute path name, otherwise the function returns `nil`.

Loading Multiple Slice Files in Ruby

You can specify as many Slice files as necessary in a single invocation of `Ice::loadSlice`, as shown below:

Ruby

```
Ice::loadSlice("Syscall.ice Process.ice")
```

Alternatively, you can call `Ice::loadSlice` several times:

Copyright 2016, ZeroC, Inc.

Ruby

```
Ice::loadSlice("Syscall.ice")
Ice::loadSlice("Process.ice")
```

If a Slice file includes another file, the default behavior of `Ice::loadSlice` generates Ruby code only for the named file. For example, suppose `Syscall.ice` includes `Process.ice` as follows:

Slice

```
// Syscall.ice
#include <Process.ice>
...
```

If you call `Ice::loadSlice("-I. Syscall.ice")`, Ruby code is not generated for the Slice definitions in `Process.ice` or for any definitions that may be included by `Process.ice`. If you also need code to be generated for included files, one solution is to load them individually in subsequent calls to `Ice::loadSlice`. However, it is much simpler, not to mention more efficient, to use the `--all` option instead:

Ruby

```
Ice::loadSlice("--all -I. Syscall.ice")
```

When you specify `--all`, `Ice::loadSlice` generates Ruby code for all Slice definitions included directly or indirectly from the named Slice files.

There is no harm in loading a Slice file multiple times, aside from the additional overhead associated with code generation. For example, this situation could arise when you need to load multiple top-level Slice files that happen to include a common subset of nested files. Suppose that we need to load both `Syscall.ice` and `Kernel.ice`, both of which include `Process.ice`. The simplest way to load both files is with a single call to `Ice::loadSlice`:

Ruby

```
Ice::loadSlice("--all -I. Syscall.ice Kernel.ice")
```

Although this invocation causes the Ice extension to generate code twice for `Process.ice`, the generated code is structured so that the interpreter ignores duplicate definitions. We could have avoided generating unnecessary code with the following sequence of steps:

Ruby

```
Ice::loadSlice("--all -I. Syscall.ice")
Ice::loadSlice("-I. Kernel.ice")
```


Copyright 2016, ZeroC, Inc.

In more complex cases, however, it can be difficult or impossible to completely avoid this situation, and the overhead of code generation is usually not significant enough to justify such an effort.

Limitations of Dynamic Code Generation in Ruby

The `Ice::loadSlice` method must be called outside of any module scope. For example, the following code is incorrect:

Ruby

```
# WRONG
module M
    Ice::loadSlice("--all -I. Syscall.ice Kernel.ice")
    ...
end
```

Static Code Generation in Ruby

You should be familiar with static code generation if you have used other Slice language mappings, such as C++ or Java. Using static code generation, the Slice compiler `slice2rb` generates Ruby code from your Slice definitions.

Compiler Output in Ruby

For each Slice file `X.ice`, `slice2rb` generates Ruby code into a file named `X.rb` in the output directory. The default output directory is the current working directory, but a different directory can be specified using the `--output-dir` option.

Include Files in Ruby

It is important to understand how `slice2rb` handles include files. In the absence of the `--all` option, the compiler does not generate Ruby code for Slice definitions in included files. Rather, the compiler translates Slice `#include` statements into Ruby `require` statements in the following manner:

1. Determine the full pathname of the included file.
2. Create the shortest possible relative pathname for the included file by iterating over each of the include directories (specified using the `-I` option) and removing the leading directory from the included file if possible.
   For example, if the full pathname of an included file is `/opt/App/slice/OS/Process.ice`, and we specified the options `-I/opt/App` and `-I/opt/App/slice`, then the shortest relative pathname is `OS/Process.ice` after removing `/opt/App/slice`.
3. Replace the `.ice` extension with `.rb`. Continuing our example from the previous step, the translated `require` statement becomes

```
require "OS/Process.rb"
```

As a result, you can use `-I` options to tailor the `require` statements generated by the compiler in order to avoid absolute pathnames and match the organizational structure of your application's source files.

Static Versus Dynamic Code Generation in Ruby

There are several issues to consider when evaluating your requirements for code generation.

Application Considerations for Code Generation in Ruby

The requirements of your application generally dictate whether you should use dynamic or static code generation. Dynamic code generation is convenient for a number of reasons:

- It avoids the intermediate compilation step required by static code generation.
- It makes the application more compact because the application requires only the Slice files, not the additional files produced by

Copyright 2016, ZeroC, Inc.

static code generation.
- It reduces complexity, which is especially helpful during testing, or when writing short or transient programs.

Static code generation, on the other hand, is appropriate in many situations:

- when an application uses a large number of Slice definitions and the startup delay must be minimized
- when it is not feasible to deploy Slice files with the application
- when a number of applications share the same Slice files
- when Ruby code is required in order to utilize third-party Ruby tools.

Mixing Static and Dynamic Generation in Ruby

You can safely use a combination of static and dynamic translation in an application. For it to work properly, you must correctly manage the include paths for Slice translation and the Ruby interpreter so that the statically-generated code can be imported properly by `require`.

For example, suppose you want to dynamically load the following Slice definitions:

Slice

```
#include <Glacier2/Session.ice>

module MyApp {
    interface MySession extends Glacier2::Session {
        // ...
    };
};
```

Whether the included file `Glacier2/Session.ice` is loaded dynamically or statically is determined by the presence of the `--all` option:

Ruby

```
sliceDir = "-I#{ENV['ICE_HOME']}/slice"

# Load Glacier2/Session.ice dynamically:
Ice::loadSlice(sliceDir + " --all MySession.ice")

# Load Glacier2/Session.ice statically:
Ice::loadSlice(sliceDir + " MySession.ice")
```

In this example, the first invocation of `loadSlice` uses the `--all` option so that code is generated dynamically for all included files. The second invocation omits `--all`, therefore the Ruby interpreter executes the equivalent of the following statement:

```
require "Glacier2/Session.rb"
```

As a result, before we can call `loadSlice` we must first ensure that the interpreter can locate the statically-generated file `Glacier2/Session.rb`. We can do this in a number of ways, including:

- adding the parent directory (e.g., `/opt/IceRuby/ruby`) to the `RUBYLIB` environment variable
- specifying the `-I` option when starting the interpreter
- modifying the search path at run time, as shown below:

Copyright 2016, ZeroC, Inc.

```
$:.unshift("/opt/IceRuby/ruby")
```

`slice2rb` Command-Line Options

The Slice-to-Ruby compiler, `slice2rb`, offers the following command-line options in addition to the standard options:

- `--all`
  Generate code for all Slice definitions, including those from included files.
- `--checksum`
  Generate checksums for Slice definitions.

See Also

- Using the Slice Compilers
- Using Slice Checksums in Ruby

Copyright 2016, ZeroC, Inc.

The main Program in Ruby

On this page:

- Initializing the Ice Run Time in Ruby
- The Ice::Application Class in Ruby
  - Catching Signals in Ruby
  - Ice::Application and Properties in Ruby
  - Limitations of Ice::Application in Ruby

Initializing the Ice Run Time in Ruby

The main entry point to the Ice run time is represented by the local interface `Ice::Communicator`. You must initialize the Ice run time by calling `Ice::initialize` before you can do anything else in your program.

`Ice::initialize` returns a reference to an instance of `Ice:::Communicator`:

Ruby

```
require 'Ice'

status = 0
ic = nil
begin
    ic = Ice::initialize(ARGV)
    # ...
rescue => ex
    puts ex
    status = 1
end

# ...
```

`Ice::initialize` accepts the argument list that is passed to the program by the operating system. The function scans the argument list for any command-line options that are relevant to the Ice run time; any such options are removed from the argument list so, when `Ice::initialize` returns, the only options and arguments remaining are those that concern your application. If anything goes wrong during initialization, `initialize` throws an exception.

Before leaving your program, you *must* call `Communicator.destroy`. The `destroy` operation is responsible for finalizing the Ice run time. In particular, `destroy` ensures that any outstanding threads are joined with and reclaims a number of operating system resources, such as file descriptors and memory. Never allow your program to terminate without calling `destroy` first; doing so has undefined behavior.

The general shape of our program is therefore as follows:

Copyright 2016, ZeroC, Inc.

Ruby

```
require 'Ice'

status = 0
ic = nil
begin
    ic = Ice::initialize(ARGV)
    # ...
rescue => ex
    puts ex
    status = 1
end

if ic
    begin
        ic.destroy()
    rescue => ex
        puts ex
        status = 1
    end
end

exit(status)
```

Note that the code places the call to `Ice::initialize` into a `begin` block and takes care to return the correct exit status to the operating system. Also note that an attempt to destroy the communicator is made only if the initialization succeeded.

The `Ice::Application` Class in Ruby

The preceding program structure is so common that Ice offers a class, `Ice::Application`, that encapsulates all the correct initialization and finalization activities. The synopsis of the class is as follows (with some detail omitted for now):

Ruby

```
module Ice
    class Application
        def main(args, configFile=nil, initData=nil)

        def run(args)

        def Application.appName()

        def Application.communicator()
    end
end
```

The intent of this class is that you specialize `Ice::Application` and implement the abstract `run` method in your derived class. Whatever code you would normally place in your main program goes into `run` instead. Using `Ice::Application`, our program looks as follows:

Copyright 2016, ZeroC, Inc.

Ruby

```
require 'Ice'

class Client < Ice::Application
    def run(args)
        # Client code here...
        return 0
    end
end

app = Client.new()
status = app.main(ARGV)
exit(status)
```

If you prefer, you can also reopen `Ice::Application` and define `run` directly:

Ruby

```
require 'Ice'

class Ice::Application
    def run(args)
        # Client code here...
        return 0
    end
end

app = Ice::Application.new()
status = app.main(ARGV)
exit(status)
```

You may also call `main` with an optional file name or an `InitializationData` structure. If you pass a configuration file name to `main`, the settings in this file are overridden by settings in a file identified by the `ICE_CONFIG` environment variable (if defined). Property settings supplied on the command line take precedence over all other settings.

The `Application.main` method does the following:

1. It installs an exception handler. If your code fails to handle an exception, `Application.main` prints the exception information before returning with a non-zero return value.
2. It initializes (by calling `Ice::initialize`) and finalizes (by calling `Communicator.destroy`) a communicator. You can get access to the communicator for your program by calling the static `communicator` accessor.
3. It scans the argument list for options that are relevant to the Ice run time and removes any such options. The argument list that is passed to your `run` method therefore is free of Ice-related options and only contains options and arguments that are specific to your application.
4. It provides the name of your application via the static `appName` method. The return value from this call is the first element of the argument vector passed to `Application.main`, so you can get at this name from anywhere in your code by calling `Ice::Application::appName` (which is often necessary for error messages).
5. It installs a signal handler that properly shuts down the communicator.

Using `Ice::Application` ensures that your program properly finalizes the Ice run time, whether your program terminates normally or in response to an exception or signal. We recommend that all your programs use this class; doing so makes your life easier. In addition `Ice::Application` also provides features for signal handling and configuration that you do not have to implement yourself when you use this

Copyright 2016, ZeroC, Inc.

class.

Catching Signals in Ruby

A program typically needs to perform some cleanup work before terminating, such as flushing database buffers or closing network connections. This is particularly important on receipt of a signal or keyboard interrupt to prevent possible corruption of database files or other persistent data.

To make it easier to deal with signals, `Ice::Application` encapsulates Ruby's signal handling capabilities, allowing you to cleanly shut down on receipt of a signal:

Ruby

```
class Application
    def Application.destroyOnInterrupt()

    def Application.ignoreInterrupt()

    def Application.callbackOnInterrupt()

    def Application.holdInterrupt()

    def Application.releaseInterrupt()

    def Application.interrupted()

    def interruptCallback(sig):
        # Default implementation does nothing.
    end
    # ...
end
```

The methods behave as follows:

- `destroyOnInterrupt`
  This method installs a signal handler that destroys the communicator if it is interrupted. This is the default behavior.
- `ignoreInterrupt`
  This method causes signals to be ignored.
- `callbackOnInterrupt`
  This function configures `Ice::Application` to invoke `interruptCallback` when a signal occurs, thereby giving the subclass responsibility for handling the signal.
- `holdInterrupt`
  This method temporarily blocks signal delivery.
- `releaseInterrupt`
  This method restores signal delivery to the previous disposition. Any signal that arrives after `holdInterrupt` was called is delivered when you call `releaseInterrupt`.
- `interrupted`
  This method returns `True` if a signal caused the communicator to shut down, `False` otherwise. This allows us to distinguish intentional shutdown from a forced shutdown that was caused by a signal. This is useful, for example, for logging purposes.
- `interruptCallback`
  A subclass implements this function to respond to signals. The function may be called concurrently with any other thread and must not raise exceptions.

By default, `Ice::Application` behaves as if `destroyOnInterrupt` was invoked, therefore our program requires no change to ensure that the program terminates cleanly on receipt of a signal. (You can disable the signal-handling functionality of `Ice::Application` by

Copyright 2016, ZeroC, Inc.

passing the constant `NoSignalHandling` to the constructor. In that case, signals retain their default behavior, that is, terminate the process.) However, we add a diagnostic to report the occurrence of a signal, so our program now looks like:

Ruby

```
require 'Ice'

class MyApplication < Ice::Application
    def run(args)
        # Client code here...

        if Ice::Application::interrupted()
            print Ice::Application::appName() + ": terminating"
        end

        return 0
    end
end

app = MyApplication.new()
status = app.main(ARGV)
exit(status)
```

`Ice::Application` and Properties in Ruby

Apart from the functionality shown in this section, `Ice::Application` also takes care of initializing the Ice run time with property values. Properties allow you to configure the run time in various ways. For example, you can use properties to control things such as the thread pool size or the trace level for diagnostic output. The `main` method of `Ice::Application` accepts an optional second parameter allowing you to specify the name of a configuration file that will be processed during initialization.

Limitations of `Ice::Application` in Ruby

`Ice::Application` is a singleton class that creates a single communicator. If you are using multiple communicators, you cannot use `Ice::Application`. Instead, you must structure your code as we saw in Hello World Application (taking care to always destroy the communicator).

See Also

- Hello World Application
- Properties and Configuration
- Communicator Initialization
- Logger Facility

Copyright 2016, ZeroC, Inc.

Using Slice Checksums in Ruby

The Slice compilers can optionally generate checksums of Slice definitions. For `slice2rb`, the `--checksum` option causes the compiler to generate code that adds checksums to the hash collection `Ice::SliceChecksums`. The checksums are installed automatically when the Ruby code is first parsed; no action is required by the application.

In order to verify a server's checksums, a client could simply compare the two hash objects using a comparison operator. However, this is not feasible if it is possible that the server might return a superset of the client's checksums. A more general solution is to iterate over the local checksums as demonstrated below:

Ruby

```
serverChecksums = ...
for i in Ice::SliceChecksums.keys
    if not serverChecksums.has_key?(i)
        # No match found for type id!
    elif Ice::SliceChecksums[i] != serverChecksums[i]
        # Checksum mismatch!
    end
end
```

In this example, the client first verifies that the server's dictionary contains an entry for each Slice type ID, and then it proceeds to compare the checksums.

See Also

- Slice Checksums

Copyright 2016, ZeroC, Inc.

Example of a File System Client in Ruby

This page presents a very simple client to access a server that implements the file system we developed in Slice for a Simple File System. The Ruby code shown here hardly differs from the code you would write for an ordinary Ruby program. This is one of the biggest advantages of using Ice: accessing a remote object is as easy as accessing an ordinary, local Ruby object. This allows you to put your effort where you should, namely, into developing your application logic instead of having to struggle with arcane networking APIs.

We now have seen enough of the client-side Ruby mapping to develop a complete client to access our remote file system. For reference, here is the Slice definition once more:

Slice

```
module Filesystem {
    interface Node {
        idempotent string name();
    };

    exception GenericError {
        string reason;
    };

    sequence<string> Lines;

    interface File extends Node {
        idempotent Lines read();
        idempotent void write(Lines text) throws GenericError;
    };

    sequence<Node*> NodeSeq;

    interface Directory extends Node {
        idempotent NodeSeq list();
    };
};
```

To exercise the file system, the client does a recursive listing of the file system, starting at the root directory. For each node in the file system, the client shows the name of the node and whether that node is a file or directory. If the node is a file, the client retrieves the contents of the file and prints them.

The body of the client code looks as follows:

Ruby

```
require 'Filesystem.rb'

# Recursively print the contents of directory "dir"
# in tree fashion. For files, show the contents of
# each file. The "depth" parameter is the current
# nesting level (for indentation).

def listRecursive(dir, depth)
    indent = ''
    depth = depth + 1
    for i in (0...depth)
```


Copyright 2016, ZeroC, Inc.

```
        indent += "\t"
    end

    contents = dir.list()

    for node in contents
        subdir = Filesystem::DirectoryPrx::checkedCast(node)
        file = Filesystem::FilePrx::uncheckedCast(node)
        print indent + node.name() + " "
        if subdir
            puts "(directory):"
            listRecursive(subdir, depth)
        else
            puts "(file):"
            text = file.read()
            for line in text
                puts indent + "\t" + line
            end
        end
    end
end

status = 0
ic = nil
begin
    # Create a communicator
    #
    ic = Ice::initialize(ARGV)

    # Create a proxy for the root directory
    #
    obj = ic.stringToProxy("RootDir:default -p 10000")

    # Down-cast the proxy to a Directory proxy
    #
    rootDir = Filesystem::DirectoryPrx::checkedCast(obj)

    # Recursively list the contents of the root directory
    #
    puts "Contents of root directory:"
    listRecursive(rootDir, 0)
rescue => ex
    puts ex
    print ex.backtrace.join("\n")
    status = 1
end

if ic
    # Clean up
    #
    begin
        ic.destroy()
    rescue => ex
        puts ex
        print ex.backtrace.join("\n")
        status = 1
```


Copyright 2016, ZeroC, Inc.

```
    end
end
```


Copyright 2016, ZeroC, Inc.

```
exit(status)
```

The program first defines the `listRecursive` function, which is a helper function to print the contents of the file system, and the main program follows. Let us look at the main program first:

1. The structure of the code follows what we saw in Hello World Application. After initializing the run time, the client creates a proxy to the root directory of the file system. For this example, we assume that the server runs on the local host and listens using the default protocol (TCP/IP) at port 10000. The object identity of the root directory is known to be `RootDir`.
2. The client down-casts the proxy to `DirectoryPrx` and passes that proxy to `listRecursive`, which prints the contents of the file system.

Most of the work happens in `listRecursive`. The function is passed a proxy to a directory to list, and an indent level. (The indent level increments with each recursive call and allows the code to print the name of each node at an indent level that corresponds to the depth of the tree at that node.) `listRecursive` calls the `list` operation on the directory and iterates over the returned sequence of nodes:

1. The code does a `checkedCast` to narrow the `Node` proxy to a `Directory` proxy, as well as an `uncheckedCast` to narrow the `Node` proxy to a `File` proxy. Exactly one of those casts will succeed, so there is no need to call `checkedCast` twice: if the `Node` *is-a* `Directory`, the code uses the `DirectoryPrx` returned by the `checkedCast`; if the `checkedCast` fails, we *know* that the Node *is-a* File and, therefore, an `uncheckedCast` is sufficient to get a `FilePrx`.
   In general, if you know that a down-cast to a specific type will succeed, it is preferable to use an `uncheckedCast` instead of a `checkedCast` because an `uncheckedCast` does not incur any network traffic.
2. The code prints the name of the file or directory and then, depending on which cast succeeded, prints `"(directory)"` or `"(file)"` following the name.
3. The code checks the type of the node:
   - If it is a directory, the code recurses, incrementing the indent level.
   - If it is a file, the code calls the `read` operation on the file to retrieve the file contents and then iterates over the returned sequence of lines, printing each line.

Assume that we have a small file system consisting of a two files and a a directory as follows:

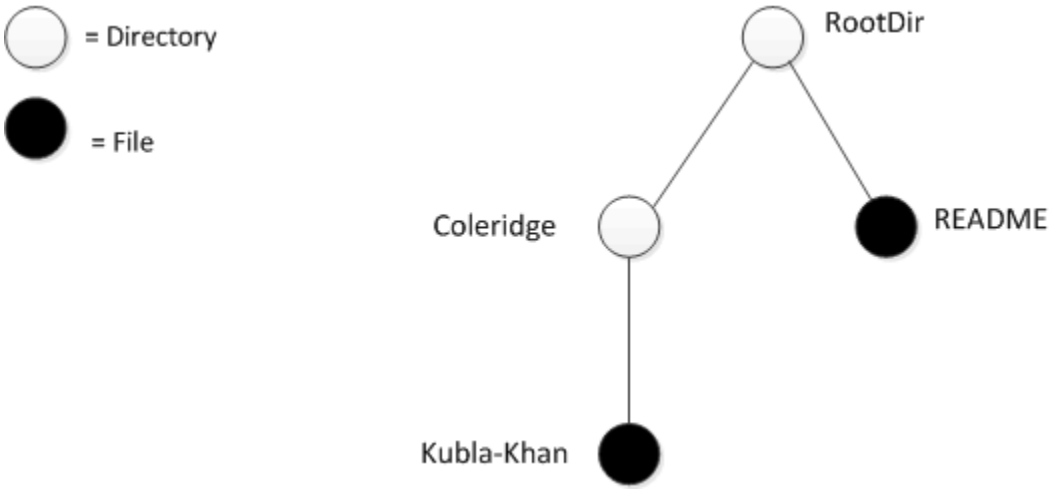

A small file system.

The output produced by the client for this file system is:

Copyright 2016, ZeroC, Inc.

```
Contents of root directory:
        README (file):
                This file system contains a collection of poetry.
        Coleridge (directory):
                Kubla_Khan (file):
                        In Xanadu did Kubla Khan
                        A stately pleasure-dome decree:
                        Where Alph, the sacred river, ran
                        Through caverns measureless to man
                        Down to a sunless sea.
```

Note that, so far, our client is not very sophisticated:

- The protocol and address information are hard-wired into the code.
- The client makes more remote procedure calls than strictly necessary; with minor redesign of the Slice definitions, many of these calls can be avoided.

We will see how to address these shortcomings in our discussions of IceGrid and object life cycle.

See Also

- Hello World Application
- Slice for a Simple File System
- Object Life Cycle
- IceGrid

Copyright 2016, ZeroC, Inc.

Properties and Configuration

Ice uses a configuration mechanism that allows you to control many aspects of the behavior of your Ice applications at run time, such as the maximum message size, the number of threads, or whether to produce network trace messages. The configuration mechanism is not only useful for configuring Ice, but also for configuring your own applications. The configuration mechanism is simple to use with a minimal API, yet flexible enough to cope with the needs of most applications.

Topics

- Properties Overview
- Configuration File Syntax
- Setting Properties on the Command Line
- Using Configuration Files
- Alternate Property Stores
- Command-Line Parsing and Initialization
- The Properties Interface
- Reading Properties
- Setting Properties
- Parsing Properties

Copyright 2016, ZeroC, Inc.

Properties Overview

Ice and its various subsystems are configured by properties. A property is a name-value pair, for example:

```
Ice.UDP.SndSize=65535
```

In this example, the *property name* is `Ice.UDP.SndSize`, and the *property value* is `65535`.

You can find a complete list of the properties used to configure Ice in the property reference.

Note that Ice reads properties that control the Ice run time and its services (that is, properties that start with one of the reserved prefixes, such as `Ice`, `Glacier2`, etc.) only once on start-up, when you create a communicator. This means that you must set Ice-related properties to their correct values *before* you create a communicator. If you change the value of an Ice-related property after that point, it is likely that the new setting will simply be ignored.

On this page:

- Property Categories
- Reserved Prefixes for Properties
- Property Name Syntax
- Property Value Syntax
- Unused Properties

Property Categories

By convention, Ice properties use the following naming scheme:

```
<application>.<category>[.<sub-category>]
```

Note that the sub-category is optional and not used by all Ice properties.

This two- or three-part naming scheme is by convention only — if you use properties to configure your own applications, you can use property names with any number of categories.

Reserved Prefixes for Properties

Ice reserves properties with the following prefixes:

- `Ice`
- `IceBox`
- `IceGrid`
- `IcePatch2`
- `IceSSL`
- `IceStorm`
- `Freeze`
- `Glacier2`

You cannot use a property beginning with one of these prefixes to configure your own application.

Property Name Syntax

A property name consists of any number of characters. For example, the following are valid property names:

Copyright 2016, ZeroC, Inc.

```
foo
Foo
foo.bar
foo bar    White space is allowed
foo=bar    Special characters are allowed
.
```

Note that there is no special significance to a period in a property name. (Periods are used to make property names more readable and are not treated specially by the property parser.)

Property names cannot contain leading or trailing white space. (If you create a property name with leading or trailing white space, that white space is silently stripped.)

Property Value Syntax

A property value consists of any number of characters. The following are examples of property values:

```
65535
yes
This is a = property value.
../../config
```

Unused Properties

During the destruction of a communicator, the Ice run time can optionally emit a warning for properties that were set but never read. To enable this warning, set `Ice.Warn.UnusedProperties` to a non-zero value. This property is useful for detecting mis-spelled properties, such as `Filesystem.MaxFilSize`. By default, the warning is disabled.

See Also

- Property Reference

Copyright 2016, ZeroC, Inc.

Configuration File Syntax

This page describes the syntax of an Ice configuration file.

On this page:

- Configuration File Format
- Special Characters in Configuration Files

Configuration File Format

A configuration file contains any number of name-value pairs, with each pair on a separate line. Empty lines and lines consisting entirely of white space characters are ignored. The # character introduces a comment that extends to the end of the current line.

Configuration files can be ASCII text files or use the UTF?8 character encoding with a byte order marker (BOM) at the beginning of the file.

Here is a simple configuration file:

```
# Example config file for Ice

Ice.MessageSizeMax = 2048    # Largest message size is 2MB
Ice.Trace.Network=3          # Highest level of tracing for network
Ice.Trace.Protocol=          # Disable protocol tracing
```

White space within property keys and values is preserved, whether escaped with a backslash or not escaped.

Leading and trailing white space is always ignored for property *names* (whether the white space is escaped or not), for example:

```
# Key white space example

Prop1     = 1           # Key is "Prop1"
 Prop2    = 2           # Key is "Prop2"
\ Prop3 \ = 3           # Key is "Prop3"
My Prop1  = 1           # Key is "My Prop1"
My\ Prop2 = 2           # Key is "My Prop2"
```

For property values, you can preserve leading and trailing white space by escaping the white space with a backslash, for example:

```
# Value white space example

My.Prop1 = a property               # Value is "a property"
My.Prop2 =     a     property       # Value is "a     property"
My.Prop3 = \ \ a     property\ \    # Value is "  a     property  "
My.Prop4 = \ \ a  \ \  property\ \  # Value is "  a     property  "
My.Prop5 = a \\ property            # Value is "a \ property"
```

This example shows that leading and trailing white space for property values is ignored unless escaped with a backslash whereas, white space that is surrounded by non-white space characters is preserved exactly, whether it is escaped or not. As usual, you can insert a literal backslash into a property value by using a double backslash.

If you set the same property more than once, the last setting prevails and overrides any previous setting. Note that assigning nothing to a

Copyright 2016, ZeroC, Inc.

property clears that property (that is, sets it to the empty string).

Ice treats properties that contain the empty string (such as `Ice.Trace.Protocol` in the preceding example) like a property that is not set at all, and we recommend that your Ice-based applications do the same. With `getProperty`, `getPropertyAsInt`, `getPropertyAsIntWithDefault`, `getPropertyAsList` and `getPropertyAsListWithDefault`, you cannot distinguish between a property that is not set and a property set to the empty string; however, `getPropertyWithDefault` allows you to make this distinction, for example:

```
// C++

// returns 3 if not set or set to the empty string
int traceProtocol = properties->getPropertyAsIntWithDefault("Ice.Trace.Protocol",
3);

// returns "3" if not set but "" if set to the empty string
string traceProtocolString =
properties->getPropertyWithDefault("Ice.Trace.Protocol", "3");
```

Property values can include characters from non-English alphabets. The Ice run time expects the configuration file to use UTF-8 encoding for such characters. (With C++, you can specify a string converter when you read the file.)

Special Characters in Configuration Files

The characters = and # have special meaning in a configuration file:

- = marks the end of the property name and the beginning of the property value
- # starts a comment that extends to the end of the line

These characters must be escaped when they appear in a property name. Consider the following examples:

```
foo\=bar=1        # Name is "foo=bar", value is "1"
foo\#bar   = 2    # Name is "foo#bar", value is "2"
foo bar  =3       # Name is "foo bar", value is "3"
```

In a property value, a # character must be escaped to prevent it from starting a comment, but an = character does not require an escape. Consider these examples:

```
A=1           # Name is "A", value is "1"
B= 2 3 4      # Name is "B", value is "2 3 4"
C=5=\#6 # 7   # Name is "C", value is "5=#6"
```

Note that, two successive backslashes in a property value become a single backslash. To get two consecutive backslashes, you must escape each one with another backslash:

```
AServer=\\\\server\dir    # Value is "\\server\dir"
BServer=\\server\\dir     # Value is "\server\dir"
```

The preceding example also illustrates that, if a backslash is not followed by a backslash, #, or =, the backslash and the character following it are both preserved.

Copyright 2016, ZeroC, Inc.

See Also

- Using Configuration Files
- Reading Properties
- Setting Properties on the Command Line
- Communicator Initialization
- C++ Strings and Character Encoding

Copyright 2016, ZeroC, Inc.

Setting Properties on the Command Line

In addition to setting properties in a configuration file, you can also set properties on the command line, for example:

```
$ ./server --Ice.UDP.SndSize=65535 --IceSSL.Trace.Security=2
```

Any command line option that begins with `--` and is followed by one of the reserved prefixes is read and converted to a property setting when you create a communicator. Property settings on the command line override settings in a configuration file. If you set the same property more than once on the same command line, the last setting overrides any previous ones.

For convenience, any property not explicitly set to a value is set to the value `1`. For example,

```
$ ./server --Ice.Trace.Protocol
```

is equivalent to

```
$ ./server --Ice.Trace.Protocol=1
```

Note that this feature only applies to properties that are set on the command line, but not to properties that are set from a configuration file.

You can also clear a property from the command line as follows:

```
$ ./server --Ice.Trace.Protocol=
```

As for properties set from a configuration file, assigning nothing to a property clears that property.

See Also

- Properties Overview
- Using Configuration Files

Copyright 2016, ZeroC, Inc.

Using Configuration Files

The ability to configure an application's properties externally provides a great deal of flexibility: you can use any combination of command-line options and configuration files to achieve the desired settings, all without having to modify your application. This page describes two ways of loading property settings from a file.

On this page:

- Prerequisites for Using Configuration Files
- The ICE_CONFIG Environment Variable
- The Ice.Config Property

Prerequisites for Using Configuration Files

The Ice run time automatically loads a configuration file during the creation of a *property set*, which is an instance of the `Ice::Properties` interface. Every communicator has its own property set from which it derives its configuration. If an application does not supply a property set when it calls `Ice::initialize` (or the equivalent in other language mappings), the Ice run time internally creates a *property set* for the new communicator.

Note however that Ice loads a configuration file automatically only when the application creates a property set using an argument vector. This occurs when the application passes an argument vector to create a property set explicitly, or when the application passes an argument vector to `Ice::initialize`.

Both of the mechanisms described below can also retrieve property settings from additional sources.

The `ICE_CONFIG` Environment Variable

Ice automatically loads the contents of the configuration file named in the `ICE_CONFIG` environment variable (assuming the prerequisites ar e met). For example:

```
$ export ICE_CONFIG=/usr/local/filesystem/config
$ ./server
```

This causes the server to read its property settings from the configuration file in `/usr/local/filesystem/config`.

If you use the `ICE_CONFIG` environment variable together with command-line options for other properties, the settings on the command line override the settings in the configuration file. For example:

```
$ export ICE_CONFIG=/usr/local/filesystem/config
$ ./server --Ice.MessageSizeMax=4096
```

This sets the value of the `Ice.MessageSizeMax` property to `4096` regardless of any setting of this property in `/usr/local/filesystem /config`.

You can use multiple configuration files by specifying a list of configuration file names separated by commas. For example:

```
$ export ICE_CONFIG=/usr/local/filesystem/config,./config
$ ./server
```

This causes property settings to be retrieved from `/usr/local/filesystem/config`, followed by any settings in the file `config` in the current directory; settings in `./config` override settings `/usr/local/filesystem/config`.

Copyright 2016, ZeroC, Inc.

The `Ice.Config` Property

The `Ice.Config` property has special meaning to the Ice run time: it determines the path name of a configuration file from which to read property settings. For example:

```
$ ./server --Ice.Config=/usr/local/filesystem/config
```

This causes property settings to be read from the configuration file in `/usr/local/filesystem/config`.

The `--Ice.Config` command-line option overrides any setting of the `ICE_CONFIG` environment variable, that is, if the `ICE_CONFIG` environment variable is set and you also use the `--Ice.Config` command-line option, the configuration file specified by the `ICE_CONFIG` environment variable is ignored.

If you use the `--Ice.Config` command-line option together with settings for other properties, the settings on the command line override the settings in the configuration file. For example:

```
$ ./server --Ice.Config=/usr/local/filesystem/config --Ice.MessageSizeMax=4096
```

This sets the value of the `Ice.MessageSizeMax` property to `4096` regardless of any setting of this property in `/usr/local/filesystem/config`. The placement of the `--Ice.Config` option on the command line has no influence on this precedence. For example, the following command is equivalent to the preceding one:

```
$ ./server --Ice.MessageSizeMax=4096 --Ice.Config=/usr/local/filesystem/config
```

Settings of the `Ice.Config` property inside a configuration file are ignored, that is, you can set `Ice.Config` only on the command line.

If you use the `--Ice.Config` option more than once, only the last setting of the option is used and the preceding ones are ignored. For example:

```
$ ./server --Ice.Config=file1 --Ice.Config=file2
```

This is equivalent to using:

```
$ ./server --Ice.Config=file2
```

You can use multiple configuration files by specifying a list of configuration file names separated by commas. For example:

```
$ ./server --Ice.Config=/usr/local/filesystem/config,./config
```

This causes property settings to be retrieved from `/usr/local/filesystem/config`, followed by any settings in the file `config` in the current directory; settings in `./config` override settings `/usr/local/filesystem/config`.

See Also

- Alternate Property Stores

Copyright 2016, ZeroC, Inc.

- The Properties Interface

Copyright 2016, ZeroC, Inc.

Alternate Property Stores

In addition to regular files, Ice also supports storing property settings in the Windows registry and Java resources.

On this page:

- Loading Properties from the Windows Registry
- Loading Properties from Java Resources

Loading Properties from the Windows Registry

You can use the Windows registry to store property settings. Property settings must be stored with a key underneath `HKEY_LOCAL_MACHINE`. To inform the Ice run time of this key, you must set the `Ice.Config` property to the key. For example:

```
$ client --Ice.Config=HKLM\MyCompany\MyApp
```

The Ice run time examines the value of `Ice.Config`; if that value begins with `HKLM`, the remainder of the property is taken to be a key to a number of string values. For the preceding example, the Ice run time looks for the key `HKEY_LOCAL_MACHINE\MyCompany\MyApp`. The string values stored under this key are used to initialize the properties.

The name of each string value is the name of the property (such as `Ice.Trace.Network`). Note that the value must be a string (even if the property setting is numeric). For example, to set `Ice.Trace.Network` to 3, you must store the string "3" as the value, not a binary or `DWORD` value.

String values in the registry can be regular strings (`REG_SZ`) or expandable strings (`REG_EXPAND_SZ`). Expandable strings allow you to include symbolic references to environment variables (such as `%ICE_HOME%`).

Depending on whether you use 32-bit or 64-bit binaries, you must set the registry keys in the corresponding 32-bit or 64-bit registry. See http://support.microsoft.com/kb/305097 for more information.

Loading Properties from Java Resources

The Ice run time for Java supports the ability to load a configuration file as a class loader resource, which is especially useful for deploying an Ice application in a self-contained JAR file. For example, suppose we define `ICE_CONFIG` as shown below:

```
$ export ICE_CONFIG=app_config
```

During the creation of a property set (which often occurs implicitly when initializing a new communicator), Ice asks the Java run time to search the application's class path for a file named `app_config`. This file might reside in the same JAR file as the application's class files, or in a different JAR file in the class path, or it might be a regular file located in one of the directories in the class path. If Java is unable to locate the configuration file in the class path, Ice attempts to open the file in the local file system.

The class path resource always takes precedence over a regular file. In other words, if a class path resource and a regular file are both present with the same path name, Ice always loads the class path resource in preference to the regular file.

The path name for a class path resource uses a relative Unix-like format such as `subdir/myfile`. Java searches for the resource relative to each JAR file or subdirectory in an application's class path.

See Also

- Using Configuration Files

Copyright 2016, ZeroC, Inc.

Command-Line Parsing and Initialization

On this page:

- Parsing Command Line Options
- The Ice.ProgramName Property

Parsing Command Line Options

When you initialize the Ice run time by calling `Ice::initialize` (C++/Ruby), `Ice.Util.initialize` (Java/C#), `Ice.initialize` (Python), or `Ice_initialize` (PHP), you can pass an argument vector to the initialization call.

For C++, `Ice::initialize` accepts a C++ *reference* to `argc`:

C++

```
namespace Ice {
    CommunicatorPtr initialize(int& argc, char* argv[]);
}
```

`Ice::initialize` parses the argument vector and initializes its property settings accordingly. In addition, it removes any arguments from `argv` that are property settings. For example, assume we invoke a server as:

```
$ ./server --myoption --Ice.Config=config -x a --Ice.Trace.Network=3 -y opt file
```

Initially, `argc` has the value 9, and `argv` has ten elements: the first nine elements contain the program name and the arguments, and the final element, `argv[argc]`, contains a null pointer (as required by the ISO C++ standard). When `Ice::initialize` returns, `argc` has the value 7 and `argv` contains the following elements:

```
./server
--myoption
-x
a
-y
opt
file
0             # Terminating null pointer
```

This means that you should initialize the Ice run time before you parse the command line for your application-specific arguments. That way, the Ice-related options are stripped from the argument vector for you so you do not need to explicitly skip them. If you use the `Ice::Application` helper class, the `run` member function is passed an argument vector with the Ice-related options already stripped. The same is true for the `runWithSession` member function called by the `Glacier2::Application` helper class.

For Java, `Ice.Util.initialize` is overloaded. The signatures are:

Copyright 2016, ZeroC, Inc.

Java

```
package Ice;
public final class Util {

    public static Communicator
    initialize();

    public static Communicator
    initialize(String[] args);

    public static Communicator
    initialize(StringSeqHolder args);

    public static Communicator
    initialize(InitializationData id);

    public static Communicator
    initialize(String[] args, InitializationData id);

    public static Communicator
    initialize(StringSeqHolder args, InitializationData id);


    // ...
}
```

The versions that accept an argument vector of type `String[]` do not strip Ice-related options for you, so, if you use one of these methods, your code must ignore options that start with one of the preserved prefixes. The versions that accept a `StringSeqHolder` behave like the C++ version and strip the Ice-related options from the passed argument vector.

In C#, the argument vector is passed by reference to the `initialize` method, allowing it to strip the Ice-related options:

Copyright 2016, ZeroC, Inc.

C#

```
namespace Ice {

    public sealed class Util {

        public static Communicator
        initialize();

        public static Communicator
        initialize(ref string[] args);

        public static Communicator
        initialize(InitializationData id);

        public static Communicator
        initialize(ref string[] args, InitializationData id);

        // ...

    }
}
```

The Python, Ruby, and PHP implementations of `initialize` have the same semantics as C++ and .NET; they expect the argument vector to be passed as a list from which all Ice-related options are removed.

If you use the `Ice.Application` helper class, the `run` method is passed the cleaned-up argument vector. The `Ice.Application` class is described separately for each language mapping.

The `Ice.ProgramName` Property

For C++, Python, and Ruby, `initialize` sets the `Ice.ProgramName` property to the name of the current program (`argv[0]`). In C#, `initialize` sets `Ice.ProgramName` to the value of `System.AppDomain.CurrentDomain.FriendlyName`. Your application code can read this property and use it for activities such as logging diagnostic or trace messages.

Even though `Ice.ProgramName` is initialized for you, you can still override its value from a configuration file or by setting the property on the command line.

For Java, the program name is not supplied as part of the argument vector — if you want to use the `Ice.ProgramName` property in your application, you must set it before initializing a communicator.

See Also

- Using Configuration Files
- Reading Properties
- Communicator Initialization
- Glacier2 Helper Classes

Copyright 2016, ZeroC, Inc.

The Properties Interface

You can use the same configuration file and command-line mechanisms to set application-specific properties. For example, we could introduce a property to control the maximum file size for our file system application:

```
# Configuration file for file system application

Filesystem.MaxFileSize=1024    # Max file size in kB
```

The Ice run time stores the `Filesystem.MaxFileSize` property like any other property and makes it accessible via the `Properties` inter face.

To access property values from within your program, you need to acquire the communicator's properties by calling `getProperties`:

Slice

```
module Ice {

    local interface Properties; // Forward declaration

    local interface Communicator {

        Properties getProperties();

        // ...
    };
};
```

The `Properties` interface is shown below:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local dictionary<string, string> PropertyDict;

    local interface Properties {

        string getProperty(string key);
        string getPropertyWithDefault(string key, string value);
        int getPropertyAsInt(string key);
        int getPropertyAsIntWithDefault(string key, int value);
        PropertyDict getPropertiesForPrefix(string prefix);

        void setProperty(string key, string value);

        StringSeq getCommandLineOptions();
        StringSeq parseCommandLineOptions(string prefix, StringSeq options);
        StringSeq parseIceCommandLineOptions(StringSeq options);

        void load(string file);

        Properties clone();
    };
};
```

Most of the operations involve reading properties, setting properties, and parsing properties.

The `Properties` interface also provides two utility operations that are useful if you need to work with multiple communicators that use different property sets:

- `clone`
  This operation makes a copy of an existing property set. The copy contains exactly the same properties and values as the original.
- `load`
  This operation accepts a path name to a configuration file and initializes the property set from that file. If the specified file cannot be read (for example, because it does not exist or the caller does not have read permission), the operation throws a `FileException`. In Java, the given path name can refer to a class loader resource or a regular file.

See Also

- Using Configuration Files
- Reading Properties
- Setting Properties
- Parsing Properties
- Alternate Property Stores

Copyright 2016, ZeroC, Inc.

Reading Properties

The `Properties` interface provides the following operations for reading property values:

- `string getProperty(string key)`
  This operation returns the value of the specified property. If the property is not set, the operation returns the empty string.
- `string getPropertyWithDefault(string key, string value)`
  This operation returns the value of the specified property. If the property is not set, the operation returns the supplied default value.
- `int getPropertyAsInt(string key)`
  This operation returns the value of the specified property as an integer. If the property is not set or contains a string that does not parse as an integer, the operation returns zero.
- `int getPropertyAsIntWithDefault(string key, int value)`
  This operation returns the value of the specified property as an integer. If the property is not set or contains a string that does not parse as an integer, the operation returns the supplied default value.
- `StringSeq getPropertyAsList(string key)`
  This operation returns the value as a list of strings. The strings must be separated by whitespace or a comma. If a string contains whitespace or a comma, it must be enclosed using single or double quotes; quotes can be themselves escaped in a quoted string, for example O'Reilly can be written as O'Reilly, "O'Reilly" or "O\'Reilly". If the property is not set or set to an empty value, the operation returns an empty list.
- `StringSeq getPropertyAsListWithDefault(string key, StringSeq value)`
  This operation returns the value as a list of strings (see `getPropertyAsList` above). If the property is not set or set to an empty value, the operation returns the supplied default value.
- `PropertyDict getPropertiesForPrefix(string prefix)`
  This operation returns all properties that begin with the specified prefix as a dictionary of type `PropertyDict`. This operation is useful if you want to extract the properties for a specific subsystem. For example,
  `getPropertiesForPrefix("Filesystem")`
  returns all properties that start with the prefix `Filesystem`, such as `Filesystem.MaxFileSize`. You can then use the usual dictionary lookup operations to extract the properties of interest from the returned dictionary.

With these operations, using application-specific properties now becomes the simple matter of initializing a communicator as usual, getting access to the communicator's properties, and examining the desired property value. For example:

C++

```
// ...

Ice::CommunicatorPtr ic;

// ...

ic = Ice::initialize(argc, argv);

// Get the maximum file size.
//
Ice::PropertiesPtr props = ic->getProperties();
Ice::Int maxSize = props->getPropertyAsIntWithDefault("Filesystem.MaxFileSize",
1024);

// ...
```

Assuming that you have created a configuration file that sets the `Filesystem.MaxFileSize` property (and that you have set the `ICE_CONFIG` variable or the `--Ice.Config` option accordingly), your application will pick up the configured value of the property.

The technique shown above allows you to obtain application-specific properties from a *configuration file*. If you also want the ability

Copyright 2016, ZeroC, Inc.

to set application-specific properties on the command line, you will need to parse command-line options for your prefix. (Calling `initialize` to create a communicator only parses those command line options having a reserved prefix.)

See Also

- The Properties Interface
- Using Configuration Files
- Setting Properties
- Parsing Properties

Copyright 2016, ZeroC, Inc.

Setting Properties

The `setProperty` operation on the `Properties` interface sets a property to the specified value:

Slice

```
module Ice {
    local interface Properties {

        void setProperty(string key, string value);

        // ...
    };
};
```

You can clear a property by setting it to the empty string.

For properties that control the Ice run time and its services (that is, properties that start with one of the reserved prefixes), this operation is useful only if you call it *before* you call `initialize`. This is because property values are usually read by the Ice run time only once, when you call `initialize`, so the Ice run time does not pay attention to a property value that is changed after you have initialized a communicator. Of course, this begs the question of how you can set a property value and have it also recognized by a communicator.

To permit you to set properties before initializing a communicator, the Ice run time provides an overloaded helper function called `createProperties` that creates a property set. In C++, the function is in the `Ice` namespace:

C++

```
namespace Ice {

PropertiesPtr createProperties(const StringConverterPtr& = 0);
PropertiesPtr createProperties(StringSeq&,
                               const PropertiesPtr& = 0,
                               const StringConverterPtr& = 0);
PropertiesPtr createProperties(int&, char*[],
                               const PropertiesPtr& = 0,
                               const StringConverterPtr& = 0);

}
```

The `StringConverter` parameter allows you to parse properties whose values contain non-ASCII characters and to correctly convert these characters into the native codeset. The converter that is passed to `createProperties` remains attached to the returned property set for the life time of the property set.

The function is overloaded to accept either an `argc`/`argv` pair or a `StringSeq`, to aid in parsing properties.

In Java, the functions are static methods of the `Util` class inside the `Ice` package:

Copyright 2016, ZeroC, Inc.

Java

```
package Ice;

public final class Util
{
    public static Properties
    createProperties();

    public static Properties
    createProperties(StringSeqHolder args);

    public static Properties
    createProperties(StringSeqHolder args, Properties defaults);

    public static Properties
    createProperties(String[] args);

    public static Properties
    createProperties(String[] args, Properties defaults);

    // ...
}
```

In C#, the `Util` class in the `Ice` namespace supplies equivalent methods:

C#

```
namespace Ice {
    public sealed class Util {
        public static Properties createProperties();
        public static Properties createProperties(ref string[] args);
        public static Properties
createProperties(ref string[] args, Properties defaults);
    }
}
```

The Python and Ruby methods reside in the `Ice` module:

```
def createProperties(args=[], defaults=None)
```

In PHP, use the `Ice_createProperties` method:

Copyright 2016, ZeroC, Inc.

PHP

```
function Ice_createProperties(args=array(), defaults=null)
```

As for `initialize`, `createProperties` strips Ice-related command-line options from the passed argument vector. (For Java, only the versions that accept a `StringSeqHolder` do this.)

The functions behave as follows:

- The parameter-less version of `createProperties` simply creates an empty property set. It does *not* check `ICE_CONFIG` for a configuration file to parse.
- The other overloads of `createProperties` accept an argument vector and a default property set. The returned property set contains all the property settings that are passed as the default, plus any property settings in the argument vector. If the argument vector sets a property that is also set in the passed default property set, the setting in the argument vector overrides the default. The overloads that accept an argument vector also look for the `--Ice.Config` option; if the argument vector specifies a configuration file, the configuration file is parsed. The order of precedence of property settings, from lowest to highest, is:
  - Property settings passed in the default parameter
  - Property settings set in the configuration file
  - Property settings in the argument vector.
    The overloads that accept an argument vector also look for the setting of the `ICE_CONFIG` environment variable and, if that variable specifies a configuration file, parse that file. (However, an explicit `--Ice.Config` option in the argument vector or the `defaults` parameter overrides any setting of the `ICE_CONFIG` environment variable.)

`createProperties` is useful if you want to ensure that a property is set to a particular value, regardless of any setting of that property in a configuration file or in the argument vector. Here is an example:

C++

```
// Get the initialized property set.
//
Ice::PropertiesPtr props = Ice::createProperties(argc, argv);

// Make sure that network and protocol tracing are off.
//
props->setProperty("Ice.Trace.Network", "0");
props->setProperty("Ice.Trace.Protocol", "0");

// Initialize a communicator with these properties.
//
Ice::InitializationData id;
id.properties = props;
Ice::CommunicatorPtr ic = Ice::initialize(id);

// ...
```

The equivalent Java code looks as follows:

Copyright 2016, ZeroC, Inc.

Java

```
Ice.StringSeqHolder argsH = new Ice.StringSeqHolder(args);
Ice.Properties properties = Ice.Util.createProperties(argsH);
properties.setProperty("Ice.Warn.Connections", "0");
properties.setProperty("Ice.Trace.Protocol", "0");
Ice.InitializationData id = new Ice.InitializationData();
id.properties = properties;
communicator = Ice.Util.initialize(id);
```

We first convert the argument array to an initialized `StringSeqHolder`. This is necessary so `createProperties` can strip Ice-specific settings. In that way, we first obtain an initialized property set, then override the settings for the two tracing properties, and then set the properties in the `InitializationData` structure.

The equivalent Python code is shown next:

Python

```
props = Ice.createProperties(sys.argv)
props.setProperty("Ice.Trace.Network", "0")
props.setProperty("Ice.Trace.Protocol", "0")
id = Ice.InitializationData()
id.properties = props
ic = Ice.initialize(id)
```

This is the equivalent code in Ruby:

Ruby

```
props = Ice::createProperties(ARGV)
props.setProperty("Ice.Trace.Network", "0")
props.setProperty("Ice.Trace.Protocol", "0")
id = Ice::InitializationData.new
id.properties = props
ic = Ice::initialize(id)
```

Finally, we present the code in PHP:

Copyright 2016, ZeroC, Inc.

PHP

```
$props = Ice_createProperties($args);
$props->setProperty("Ice.Trace.Network", "0");
$props->setProperty("Ice.Trace.Protocol", "0");
$id = new Ice_InitializationData();
$id->properties = $props;
$ic = Ice_initialize($id);
```

See Also

- The Properties Interface
- Using Configuration Files
- Command-Line Parsing and Initialization
- Reading Properties
- Parsing Properties
- Communicator Initialization
- C++ Strings and Character Encoding

Copyright 2016, ZeroC, Inc.

Parsing Properties

The `Properties` interface provides several operations for converting properties to and from command-line options.

On this page:

- Converting Properties to Command-Line Options
- Converting Command-Line Options to Properties
- Converting Reserved Command-Line Options to Properties

Converting Properties to Command-Line Options

The `getCommandLineOptions` operation converts an initialized set of properties into a sequence of equivalent command-line options:

Slice

```
module Ice {
    local interface Properties {

        StringSeq getCommandLineOptions();

        // ...
    };
};
```

For example, if you have set the `Filesystem.MaxFileSize` property to 1024 and call `getCommandLineOptions`, the setting is returned as the string `"--Filesystem.MaxFileSize=1024"`. This operation is useful for diagnostic purposes, for example, to dump the setting of all properties to a logging facility, or if you want to fork a new process with the same property settings as the current process.

Converting Command-Line Options to Properties

The `parseCommandLineOptions` operation examines the passed argument vector for command-line options that have the specified prefix:

Slice

```
module Ice {
    local interface Properties {

        StringSeq parseCommandLineOptions(string prefix, StringSeq options);

        // ...
    };
};
```

All options having the form `--`*`prefix`*`.Key=Value` are converted to property settings (that is, they initialize the corresponding properties). The operation returns an argument vector containing the options that did not match the prefix.

The value for *`prefix`* has an implicit trailing period if one is not present. For example, when calling `parseCommandLineOptions` with a prefix value of `"File"`, the option `--File.Owner=root` would match but the option `--Filesystem.MaxFileSize=1024` would not match.

The operation parses command-line options using the same syntax rules as for properties in a configuration file. However, the user's command shell can cause differences in parsing behavior. Suppose we define the following property in a configuration file:

Copyright 2016, ZeroC, Inc.

```
MyApp.Home=C:\Program Files\MyApp
```

The presence of whitespace in the property definition is not an issue in a configuration file but can be an issue on the command line, where the equivalent option is `--MyApp.Home=C:\Program Files\MyApp`. If the user is not careful, the program may receive this as two separate options: `--MyApp.Home=C:\Program` and `Files\MyApp`. In the end, it is the user's responsibility to ensure that the property's complete key and value are contained within a single command-line option.

Because `parseCommandLineOptions` expects a sequence of strings, but C++ programs are used to dealing with `argc` and `argv`, Ice provides two utility functions that convert an `argc`/`argv` vector into a sequence of strings and vice-versa:

C++

```
namespace Ice {

    StringSeq argsToStringSeq(int argc, char* argv[]);

    void stringSeqToArgs(const StringSeq& args, int& argc, char* argv[]);

}
```

You need to use `parseCommandLineOptions` (and the utility functions) if you want to permit application-specific properties to be set from the command line. For example, to allow the `--Filesystem.MaxFileSize` option to be used on the command line, we need to initialize our program as follows:

Copyright 2016, ZeroC, Inc.

C++

```
int
main(int argc, char* argv[])
{
    // Create an empty property set.
    //
    Ice::PropertiesPtr props = Ice::createProperties();

    // Convert argc/argv to a string sequence.
    //
    Ice::StringSeq args = Ice::argsToStringSeq(argc, argv);

    // Strip out all options beginning with --Filesystem.
    //
    args = props->parseCommandLineOptions("Filesystem", args);

    // args now contains only those options that were not
    // stripped. Any options beginning with --Filesystem have
    // been converted to properties.

    // Convert remaining arguments back to argc/argv vector.
    //
    Ice::stringSeqToArgs(args, argc, argv);

    // Initialize communicator.
    //
    Ice::InitializationData id;
    id.properties = props;
    Ice::CommunicatorPtr ic = Ice::initialize(argc, argv, id);

    // At this point, argc/argv only contain options that
    // set neither an Ice property nor a Filesystem property,
    // so we can parse these options as usual.
    //
    // ...
}
```

Using this code, any options beginning with `--Filesystem` are converted to properties and are available via the property lookup operations as usual. The call to `initialize` then removes any Ice-specific command-line options so, once the communicator is created, `argc`/`argv` only contains options and arguments that are not related to setting either a filesystem or an Ice property.

An easier way to achieve the same thing is to use the overload of `Ice::initialize` that accepts a string sequence, instead of an `argc`/`argv` pair:

Copyright 2016, ZeroC, Inc.

C++

```
int
main(int argc, char* argv[])
{
    // Create an empty property set.
    //
    Ice::PropertiesPtr props = Ice::createProperties();

    // Convert argc/argv to a string sequence.
    //
    Ice::StringSeq args = Ice::argsToStringSeq(argc, argv);

    // Strip out all options beginning with --Filesystem.
    //
    args = props->parseCommandLineOptions("Filesystem", args);

    // args now contains only those options that were not
    // stripped. Any options beginning with --Filesystem have
    // been converted to properties.

    // Initialize communicator.
    //
    Ice::InitializationData id;
    id.properties = props;
    Ice::CommunicatorPtr ic = Ice::initialize(args, id);

    // At this point, args only contains options that
    // set neither an Ice property nor a Filesystem property,
    // so we can parse these options as usual.
    //
    // ...
}
```

This version of the code avoids having to convert the string sequence back into an `argc`/`argv` pair before calling `Ice::initialize`.

Converting Reserved Command-Line Options to Properties

The `parseIceCommandLineOptions` operation behaves like `parseCommandLineOptions`, but removes the reserved Ice-specific options from the argument vector:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local interface Properties {

        StringSeq parseIceCommandLineOptions(StringSeq options);

        // ...
    };
};
```

This operation is also used internally by the Ice run time to parse Ice-specific options in `initialize`.

See Also

- Properties Overview
- The Properties Interface
- Reading Properties
- Setting Properties
- Command-Line Parsing and Initialization
- Logger Facility

Copyright 2016, ZeroC, Inc.

Communicator and other Core Local Features

This section presents several core features of the Ice run-time, that are purely local to your client or server application.

Topics

- Communicators
- Communicator Initialization
- Object Identity
- Plug-in Facility

Copyright 2016, ZeroC, Inc.

Communicators

The main entry point to the Ice run time is represented by the local interface `Ice::Communicator`. An instance of `Ice::Communicator` is associated with a number of run-time resources:

- Client-side thread pool
  The client-side thread pool is used to process replies to asynchronous method invocations (AMI), to avoid deadlocks in callbacks, and to process incoming requests on bidirectional connections.
- Server-side thread pool
  Threads in this pool accept incoming connections and handle requests from clients.
- Configuration properties
  Various aspects of the Ice run time can be configured via properties. Each communicator has its own set of such configuration properties.
- Object factories
  In order to instantiate classes that are derived from a known base type, the communicator maintains a set of object factories that can instantiate the class on behalf of the Ice run time. Object factories are discussed in each client-side language mapping.
- Logger object
  A logger object implements the `Ice::Logger` interface and determines how log messages that are produced by the Ice run time are handled.
- Default router
  A router implements the `Ice::Router` interface. Routers are used by Glacier2 to implement the firewall functionality of Ice.
- Default locator
  A locator is an object that resolves an object identity to a proxy. A locator object is implemented by a location service.
- Plug-in manager
  Plug-ins are objects that add features to a communicator. For example, IceSSL is implemented as a plug-in. Each communicator has a plug-in manager that implements the `Ice::PluginManager` interface and provides access to the set of plug-ins for a communicator.
- Object adapters
  Object adapters dispatch incoming requests and take care of passing each request to the correct servant.

Object adapters and objects that use different communicators are completely independent from each other. Specifically:

- Each communicator uses its own thread pool. This means that if, for example, one communicator runs out of threads for incoming requests, only objects using that communicator are affected. Objects using other communicators have their own thread pool and are therefore unaffected.
- Collocated invocations across different communicators are not optimized, whereas collocated invocations using the same communicator bypass much of the overhead of call dispatch.

Typically, servers use only a single communicator but, occasionally, multiple communicators can be useful. For example, IceBox, uses a separate communicator for each Ice service it loads to ensure that different services cannot interfere with each other. Multiple communicators are also useful to avoid thread starvation: if one service runs out of threads, this leaves the remaining services unaffected.

The communicator's interface is defined in Slice. Part of this interface looks as follows:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local interface Communicator {
        string proxyToString(Object* obj);
        Object* stringToProxy(string str);
        PropertyDict proxyToProperty(Object* proxy, string property);
        Object* propertyToProxy(string property);
        Identity stringToIdentity(string str);
        string identityToString(Identity id);
        ObjectAdapter createObjectAdapter(string name);
        ObjectAdapter createObjectAdapterWithEndpoints(string name,
string endpoints);
        ObjectAdapter createObjectAdapterWithRouter(string name, Router* rtr);
        void shutdown();
        void waitForShutdown();
        bool isShutdown();
        void destroy();
        // ...
    };
    // ...
};
```

The communicator offers a number of operations:

- `proxyToString`
  `stringToProxy`

  These operations allow you to convert a proxy into its stringified representation and vice versa. Instead of calling `proxyToString` on the communicator, you can also use the `ice_toString` proxy method to stringify it. However, you can only stringify non-null proxies that way — to stringify a null proxy, you must use `proxyToString`. (The stringified representation of a null proxy is the empty string.)

- `proxyToProperty`
  `propertyToProxy`

  `proxyToProperty` returns the set of proxy properties for the supplied proxy. The `property` parameter specifies the base name for the properties in the returned set. `propertyToProxy` retrieves the configuration property with the given name and converts its value into a proxy. A null proxy is returned if no property is found with the specified name.

- `identityToString`
  `stringToIdentity`

  These operations allow you to convert an identity to a string and vice versa.

- `createObjectAdapter`
  `createObjectAdapterWithEndpoints`
  `createObjectAdapterWithRouter`

  These operations create a new object adapter. Each object adapter is associated with zero or more transport endpoints. Typically, an object adapter has a single transport endpoint. However, an object adapter can also offer multiple endpoints. If so, these endpoints each lead to the same set of objects and represent alternative means of accessing these objects. This is useful, for example, if a server is behind a firewall but must offer access to its objects to both internal and external clients; by binding the adapter to both the internal and external interfaces, the objects implemented in the server can be accessed via either interface.

  An object adapter also can have no endpoint at all. In that case, the adapter can only be reached via collocated invocations originating from the same communicator as is used by the adapter.

Copyright 2016, ZeroC, Inc.

Whereas `createObjectAdapter` determines its transport endpoints from configuration information, `createObjectAdapterWithEndpoints` allows you to supply the transport endpoints for the new adapter. Typically, you should use `createObjectAdapter` in preference to `createObjectAdapterWithEndpoints`. Doing so keeps transport-specific information, such as host names and port numbers, out of the source code and allows you to reconfigure the application by changing a property (and so avoid recompilation when a transport endpoint needs to be changed).

`createObjectAdapterWithRouter` creates a routed object adapter that allows clients to receive callbacks from servers that are behind a router.

The newly-created adapter uses its name as a prefix for a collection of configuration properties that tailor the adapter's behavior. By default, the adapter prints a warning if other properties are defined having the same prefix, but you can disable this warning using the property `Ice.Warn.UnknownProperties`.

- `shutdown`

  This operation shuts down the server side of the Ice run time:
  - Operation invocations that are in progress at the time `shutdown` is called are allowed to complete normally. `shutdown` does *not* wait for these operations to complete; when `shutdown` returns, you know that no new incoming requests will be dispatched, but operations that were already in progress at the time you called `shutdown` may still be running. You can wait for still-executing operations to complete by calling `waitForShutdown`.
  - Operation invocations that arrive after the server has called `shutdown` either fail with a `ConnectFailedException` or are transparently redirected to a new instance of the server (via IceGrid).
  - Note that `shutdown` initiates deactivation of all object adapters associated with the communicator, so attempts to use an adapter once `shutdown` has completed raise an `ObjectAdapterDeactivatedException`.

- `waitForShutdown`

  On the server side, this operation suspends the calling thread until the communicator has shut down (that is, until no more operations are executing in the server). This allows you to wait until the server is idle before you destroy the communicator.

  On the client side, `waitForShutdown` simply waits until another thread has called `shutdown` or `destroy`.

- `isShutdown`

  This operation returns true if `shutdown` has been invoked on the communicator. A return value of true does not necessarily indicate that the shutdown process has completed, only that it has been initiated. An application that needs to know whether shutdown is complete can call `waitForShutdown`. If the blocking nature of `waitForShutdown` is undesirable, the application can invoke it from a separate thread.

- `destroy`

  This operation destroys the communicator and all its associated resources, such as threads, communication endpoints, object adapters, and memory resources. Once you have destroyed the communicator (and therefore destroyed the run time for that communicator), you must not call any other Ice operation (other than to create another communicator).

  It is imperative that you call `destroy` before you leave the `main` function of your program. Failure to do so results in undefined behavior.

  Calling `destroy` before leaving `main` is necessary because `destroy` waits for all running threads to terminate before it returns. If you leave `main` without calling `destroy`, you will leave `main` with other threads still running; many threading packages do not allow you to do this and end up crashing your program.

  If you call `destroy` without calling `shutdown`, the call waits for all executing operation invocations to complete before it returns (that is, the implementation of `destroy` implicitly calls `shutdown` followed by `waitForShutdown`). `shutdown` (and, therefore, `destroy`) deactivates all object adapters that are associated with the communicator. Since `destroy` blocks until all operation invocations complete, a servant will deadlock if it invokes `destroy` on its own communicator while executing a dispatched operation.

  On the client side, calling `destroy` while operations are still executing causes those operations to terminate with a `CommunicatorDestroyedException`.

See Also

- Properties and Configuration
- Object Adapters
- Object Adapter Endpoints
- Object Identity

Copyright 2016, ZeroC, Inc.

- Obtaining Proxies
- Proxy Methods
- Logger Facility
- The Ice Threading Model
- Bidirectional Connections
- Locators
- IceSSL
- Glacier2
- IceBox

Copyright 2016, ZeroC, Inc.

Communicator Initialization

During the creation of a communicator, the Ice run time initializes a number of features that affect the communicator's operation. Once set, these features remain in effect for the life time of the communicator, that is, you cannot change these features after you have created a communicator. Therefore, if you want to customize these features, you must do so when you create the communicator.

The following features can be customized at communicator creation time:

- the property set
- the logger object
- the instrumentation observer
- the narrow and wide string converters (C++ only)
- the thread notification hook
- the dispatcher
- the compact ID resolver (used by the streaming interfaces when extracting Ice objects)
- the class loader (Java only)
- the batch request interceptor

To establish these features, you initialize a structure or class of type `InitializationData` with the relevant settings. For C++ the structure is defined as follows:

C++

```
namespace Ice {
    struct InitializationData {
        PropertiesPtr properties;
        LoggerPtr logger;
        Instrumentation::CommunicatorObserverPtr observer;
        StringConverterPtr stringConverter;
        WstringConverterPtr wstringConverter;
        ThreadNotificationPtr threadHook;
        DispatcherPtr dispatcher;
        CompactIdResolverPtr compactIdResolver;
        BatchRequestInterceptorPtr batchRequestInterceptor;
    };
}
```

For languages other than C++, `InitializationData` is a class with all data members public. (The data members supported by this class vary with each language mapping.)

For C++, `Ice::initialize` is overloaded as follows:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {
    CommunicatorPtr initialize(int&, char*[],
                const InitializationData& = InitializationData(),
                Int = ICE_INT_VERSION);
    CommunicatorPtr initialize(StringSeq&,
                const InitializationData& = InitializationData(),
                Int = ICE_INT_VERSION);
    CommunicatorPtr initialize(
                const InitializationData& = InitializationData()
                Int = ICE_INT_VERSION);
}
```

The versions of `initialize` that accept an argument vector look for Ice-specific command-line options and remove them from the argument vector, as described in the C++ language mapping. The version without an `argc`/`argv` pair is useful if you want to prevent property settings for a program from being changed by command-line arguments — you can use the `Properties` interface for this purpose.

To set a feature, you set the corresponding field in the `InitializationData` structure and pass the structure to `initialize`. For example, to establish a custom logger of type `MyLogger`, you can use:

```
Ice::InitializationData id;
id.logger = new MyLoggerI;
Ice::CommunicatorPtr ic = Ice::initialize(argc, argv, id);
```

For Java, C#, and Objective-C, `Ice.Util.initialize` is overloaded similarly (as is `Ice.initialize` for Python, `Ice::initialize` f or Ruby, and `Ice_initialize` for PHP), so you can pass an `InitializationData` instance either with or without an argument vector.

For Objective-C, the method name is `createCommunicator`.

Note that you must supply an argument vector if you want `initialize` to look for a configuration file in the `ICE_CONFIG` environment variable.

See Also

- The Server-Side main Function in C++
- The Server-Side main Function in Objective-C
- Command-Line Parsing and Initialization
- The Properties Interface

Copyright 2016, ZeroC, Inc.

Object Identity

On this page:

- The Ice::Identity Type
- Syntax for Stringified Identities
- Identity Helper Functions

The `Ice::Identity` Type

Each Ice object has an object identity defined as follows:

Slice

```
module Ice {
    struct Identity {
        string name;
        string category;
    };
};
```

As you can see, an object identity consists of a pair of strings, a `name` and a `category`. The complete object identity is the combination of n ame and `category`, that is, for two identities to be equal, both `name` and `category` must be the same. The `category` member is usually the empty string, unless you are using servant locators, default servants or callbacks with Glacier2.

If `name` is an empty string, `category` must be the empty string as well. (An identity with an empty `name` and a non-empty `category` is illegal.) If a proxy contains an identity in which `name` is empty, Ice interprets that proxy as a null proxy.

Object identities can be represented as strings; the category part appears first and is followed by the name; the two components are separated by a `/` character, for example:

```
Factory/File
```

In this example, `Factory` is the category, and `File` is the name. If the `name` or `category` member themselves contain a `/` character, the stringified representation escapes the `/` character with a `\`, for example:

```
Factories\/Factory/Node\/File
```

In this example, the category is `Factories/Factory` and the name is `Node/File`.

Syntax for Stringified Identities

You rarely need to write identities as strings because, typically, your code will be using the identity helper functions `identityToString` an d `stringToIdentity`, or simply deal with proxies instead of identities. However, on occasion, you will need to use stringified identities in configuration files. If the identities happen to contain meta-characters (such as a slash or backslash), or characters outside the printable ASCII range, these characters must be escaped in the stringified representation. Here are rules that the Ice run time applies when parsing a stringified identity:

1. The parser scans the stringified identity for an unescaped slash character (`/`). If such a slash character can be found, the substrings to the left and right of the slash are parsed as the `category` and `name` members of the identity, respectively; if no such slash character can be found, the entire string is parsed as the `name` member of the identity, and the `category` member is the empty string.
2. Each of the `category` (if present) and `name` substrings is parsed according to the following rules:
    - All characters in the string must be in the ASCII range 32 (space) to 126 (~); characters outside this range cause the parse

Copyright 2016, ZeroC, Inc.

to fail.
- Any character that is not part of an escape sequence is treated as that character.
- The parser recognizes the following escape sequences and replaces them with their equivalent character:
  `\\` (backslash)
  `\'` (single quote)
  `\"` (double quote)
  `\b` (space)
  `\f` (form feed)
  `\n` (new line)
  `\r` (carriage return)
  `\t` (tab)
- An escape sequence of the form `\`*o*, `\`*oo*, or `\`*ooo* (where *o* is a digit in the range 0 to 7) is replaced with the ASCII character with the corresponding octal value. Parsing for octal digits allows for at most three consecutive digits, so the string `\0763` is interpreted as the character with octal value 76 (>) followed by the character 3. Parsing for octal digits terminates as soon as it encounters a character that is not in the range 0 to 7, so `\7x` is the character with octal value 7 (bell) followed by the character `x`. Octal escape sequences must be in the range 0 to 255 (octal 000 to 377); escape sequences outside this range cause a parsing error. For example, `\539` is an illegal escape sequence.
- If a character follows a backslash, but is not part of a recognized escape sequence, the backslash is ignored, so `\x` is the character `x`.

Identity Helper Functions

To make conversion of identities to and from strings easier, Ice provides functions to convert an `Identity` to and from a native string, using the string format described in the preceding paragraph.

In C++, these helper functions are in the Ice namespace:

C++

```
namespace Ice
{
    std::string identityToString(const Identity&);
    Identity    stringToIdentity(const std::string&);
}
```

For Java, the utility functions are in the `Ice.Util` class and are defined as:

Java

```
package Ice;

public final class Util {
    public static String   identityToString(Identity id);
    public static Identity stringToIdentity(String s);
}
```

For C#, the utility functions are in the `Ice.Util` class and are defined as:

Copyright 2016, ZeroC, Inc.

C#

```
namespace Ice
{
    public sealed class Util
    {
        public static string   identityToString(Identity id);
        public static Identity stringToIdentity(string s);
    }
}
```

The Python functions are in the `Ice` module:

Python

```
def identityToString(ident)
def stringToIdentity(str)
```

The PHP functions are in the `Ice` namespace:

PHP

```
function identityToString($str)
function stringToIdentity($ident)
```

The Ruby functions are in the `Ice` module:

Ruby

```
def Ice.identityToString(str)
def Ice.stringToIdentity(ident)
```

The local interface Communicator provides two equivalent operations, `identityToString` and `stringToIdentity`:

Slice

```
module Ice
{
local interface Communicator {
   Identity stringToIdentity(string str);
   string   identityToString(Identity ident);
};
};
```


Copyright 2016, ZeroC, Inc.

In Ice 3.6.1 and prior releases, Ice provides these helper functions in C++, Objective-C, Ruby, PHP and Python only on the Communicator interface; the freestanding functions were accidentally left out.

See Also

- Servant Activation and Deactivation
- Servant Locators
- Default Servants
- C++ Strings and Character Encoding
- Glacier2

Copyright 2016, ZeroC, Inc.

Plug-in Facility

Ice provides a plug-in facility that allows you to load new features into your application, without changing this application's code. Depending on the programming language, a plug-in is packaged as a shared library or a static library, a set of Java classes, or a .NET assembly.

You control the loading and initialization of plug-ins with Ice configuration properties, and the plug-ins themselves are configured through Ice configuration properties.

Ice relies on plug-ins to implement a number of features, such as IceSSL (a secure transport for Ice over SSL/TLS).

This section describes the plug-in facility in more detail and demonstrates how to implement an Ice plug-in.

Topics

- Plug-in API
- Plug-in Configuration
- Advanced Plug-in Topics

See Also

- IceSSL
- Logger Plug-ins
- The Ice String Converter Plug-in

Copyright 2016, ZeroC, Inc.

Plug-in API

On this page:

- The Plugin Interface
- C++ Plug-in Factory
- Java Plug-in Factory
- C# Plug-in Factory

The `Plugin` Interface

The plug-in facility defines a local Slice interface that all plug-ins must implement:

Slice

```
module Ice {
local interface Plugin {
    void initialize();
    void destroy();
};
};
```

The lifecycle of an Ice plug-in is structured to accommodate dependencies between plug-ins, such as when a logger plug-in needs to use IceSSL for its logging activities. Consequently, a plug-in object's lifecycle consists of four phases:

- Construction
  The Ice run time uses a language-specific factory API for instantiating plug-ins. During construction, a plug-in can acquire resources but must not spawn new threads or perform activities that depend on other plug-ins.
- Initialization
  After all plug-ins have been constructed, the Ice run time invokes `initialize` on each plug-in. The order in which plug-ins are initialized is undefined by default but can be customized using a configuration property. If a plug-in has a dependency on another plug-in, you must configure the Ice run time so that initialization occurs in the proper order. In this phase it is safe for a plug-in to spawn new threads; it is also safe for a plug-in to interact with other plug-ins and use their services, as long as those plug-ins have already been initialized. If `initialize` raises an exception, the Ice run time invokes `destroy` on all plug-ins that were successfully initialized (in the reverse order of initialization) and raises the original exception to the application.
- Active
  The active phase spans the time between initialization and destruction. Plug-ins must be designed to operate safely in the context of multiple threads.
- Destruction
  The Ice run time invokes `destroy` on each plug-in in the reverse order of initialization.

  ⚠ A plug-in's `destroy` implementation must not make remote invocations.

This lifecycle is repeated for each new communicator that an application creates and destroys.

C++ Plug-in Factory

In C++, a plug-in factory is a function with C linkage and the following signature:

Copyright 2016, ZeroC, Inc.

C++

```
extern "C"
{
Ice::Plugin*
functionName(const Ice::CommunicatorPtr& communicator,
             const std::string& name,
             const Ice::StringSeq& args)
{
   ...
}
}
```

You can define the function with any name you wish.

Since the function uses C linkage, it must return the plug-in object as a regular C++ pointer and not as an Ice smart pointer. Furthermore, the function must not raise C++ exceptions; if an error occurs, the function must return zero. The arguments to the function consist of the communicator that is in the process of being initialized, the name assigned to the plug-in, and any arguments that were specified in the plug-in's configuration.

If your plug-in and the associated factory function are packaged in a shared library or DLL loaded at run time, you need to export this function from the shared library or DLL. We provide the macro `ICE_DECLSPEC_EXPORT` for this purpose:

C++

```
#if defined(_WIN32)
#   define ICE_DECLSPEC_EXPORT __declspec(dllexport)
#elif defined(__GNUC__)
#   define ICE_DECLSPEC_EXPORT __attribute__((visibility ("default")))
#elif defined(__SUNPRO_CC)
#   define ICE_DECLSPEC_EXPORT __global
...
```

Simply add `ICE_DECLSPEC_EXPORT` to the definition of your plug-in factory:

C++

```
extern "C"
{
ICE_DECLSPEC_EXPORT Ice::Plugin*
functionName(const Ice::CommunicatorPtr& communicator,
             const std::string& name,
             const Ice::StringSeq& args)
{
   ...
}
}
```

If your application doesn't want to rely on the dynamic loading of your plug-in shared library or DLL at run time, or if your plug-in is packaged in a static library, it's also possible to link the plug-in into your application and call `Ice::registerPluginFactory` in your main application's code to register the plug-in before you initialize Ice communicators. For example:

Copyright 2016, ZeroC, Inc.

C++

```
MyApp::MyApp()
{
   // Load/link the "IceSSL" plug-in before communicator initialization

   Ice::registerPluginFactory("IceSSL", createIceSSL, false);
}
```

The `registerPluginFactory` function registers the plug-in's factory function with the Ice run time. It returns `void`, and accepts the following parameters:

- `const string&`
  The name of the plug-in.
- `PLUGIN_FACTORY`
  A pointer to the plug-in factory function.
- `bool`
  When `true`, the plug-in is always loaded (created) during communicator initialization, even if `Ice.Plugin.`*`name`* is not set. When `false`, the plug-in is loaded (created) during communication initialization only if `Ice.Plugin.`*`name`* is set to a non-empty value (e.g.: `Ice.Plugin.IceSSL=1`).

Java Plug-in Factory

In Java, a plug-in factory must implement the `Ice.PluginFactory` interface:

Java

```
package Ice;

public interface PluginFactory {
    Plugin create(Communicator communicator, String name, String[] args);
}
```

The arguments to the `create` method consist of the communicator that is in the process of being initialized, the name assigned to the plug-in, and any arguments that were specified in the plug-in's configuration.

The `create` method can return `null` to indicate that a general error occurred, or it can raise `PluginInitializationException` to provide more detailed information. If any other exception is raised, the Ice run time wraps it inside an instance of `PluginInitialization Exception`.

C# Plug-in Factory

In .NET, a plug-in factory must implement the `Ice.PluginFactory` interface:

Copyright 2016, ZeroC, Inc.

C#

```
namespace Ice {
    public interface PluginFactory
    {
        Plugin create(Communicator communicator, string name, string[] args);
    }
}
```

The arguments to the `create` method consist of the communicator that is in the process of being initialized, the name assigned to the plug-in, and any arguments that were specified in the plug-in's configuration.

The `create` method can return `null` to indicate that a general error occurred, or it can raise `PluginInitializationException` to provide more detailed information. If any other exception is raised, the Ice run time wraps it inside an instance of `PluginInitialization Exception`.

See Also

- Plug-in Configuration
- Advanced Plug-in Topics

Copyright 2016, ZeroC, Inc.

Plug-in Configuration

Plug-ins are installed using a configuration property of the following form:

```
Ice.Plugin.Name=entry_point [arg ...]
```

In most cases you can assign an arbitrary name to a plug-in. In the case of IceSSL, however, the plug-in requires that its name be `IceSSL`.

The value of *`entry_point`* is a language-specific representation of the plug-in's factory. In C++, it consists of the path name of the shared library or DLL containing the factory function, along with the name of the factory function. In Java and .NET, the entry point specifies the factory class.

The language-specific nature of plug-in properties can present a problem when applications that are written in multiple implementation languages attempt to share a configuration file. Ice supports an alternate syntax for plug-in properties that alleviates this issue:

```
Ice.Plugin.name.cpp=...    # C++ plug-in
Ice.Plugin.name.java=...   # Java plug-in
Ice.Plugin.name.clr=...    # .NET (Common Language Runtime) plug-in
```

Plug-in properties having a suffix of `.cpp`, `.java`, or `.clr` are loaded only by the appropriate Ice run time and ignored by others.

After extracting the plug-in's entry point from the property value, any remaining text is parsed using semantics similar to that of command-line arguments. Whitespace separates the arguments, and any arguments that contain whitespace must be enclosed in quotes:

```
Ice.Plugin.MyPlugin=entry_point --load "C:\Data Files\config.dat"
```

Ice passes these arguments to the plug-in's entry point during loading.

See Also

- Ice.Plugin.*
- Ice.InitPlugins
- Ice.PluginLoadOrder
- IceSSL

Copyright 2016, ZeroC, Inc.

Advanced Plug-in Topics

This page discusses additional aspects of the Ice plug-in facility that may be of use to applications with special requirements.

On this page:

- Plug-in Dependencies
- The Plug-in Manager
- Delayed Plug-in Initialization

Plug-in Dependencies

If a plug-in has a dependency on another plug-in, you must ensure that Ice initializes the plug-ins in the proper order. Suppose that a custom plug-in depends on IceSSL; for example, it may need to make secure invocations on another server. We start with the following C++ configuration:

```
Ice.Plugin.IceSSL=IceSSL:createIceSSL
Ice.Plugin.MyPlugin=MyPlugin:createMyPlugin
```

The problem with this configuration is that it does not specify the order in which the plug-ins should be loaded and initialized. If the Ice run time happens to initialize `MyPlugin` first, the plug-in's `initialize` method will fail if it attempts to use the services of the uninitialized IceSSL plug-in.

To remedy the situation, we need to add one more property:

```
Ice.Plugin.IceSSL=IceSSL:createIceSSL
Ice.Plugin.MyPlugin=MyPlugin:createMyPlugin
Ice.PluginLoadOrder=IceSSL, MyPlugin
```

Using the `Ice.PluginLoadOrder` property we can guarantee that the plug-ins are loaded in the correct order.

> Plug-ins added manually via the plug-in manager are appended to the end of the plug-in list, in order of addition. The last plug-in added is the first to be destroyed.

The Plug-in Manager

`PluginManager` is the name of an internal Ice object that is responsible for managing all aspects of Ice plug-ins. This object supports a local Slice interface of the same name, and an application can obtain a reference to this object using the following communicator operation:

Slice

```
module Ice {
local interface Communicator {
    PluginManager getPluginManager();
    // ...
};
};
```

The `PluginManager` interface offers three operations:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
local interface PluginManager {
    void initializePlugins();
    Plugin getPlugin(string name);
    void addPlugin(string name, Plugin pi);
};
};
```

The `initializePlugins` operation is used in special cases when an application needs to manually initialize one or more plug-ins, as discussed in the next section.

The `getPlugin` operation returns a reference to a specific plug-in. The `name` argument must match an installed plug-in, otherwise the operation raises `NotRegisteredException`. This operation is useful when a plug-in exports an interface that an application can use to query or customize its attributes or behavior.

Finally, `addPlugin` provides a way for an application to install a plug-in directly, without the use of a configuration property. This plug-in's `initialize` operation will be invoked if `initializePlugins` has not yet been called on the plug-in manager. If `initializePlugins` has already been called before a plug-in is added, Ice does not invoke `initialize` on the plug-in, but does invoke `destroy` during communicator destruction.

Delayed Plug-in Initialization

It is sometimes necessary for an application to manually configure a plug-in prior to its initialization. For example, SSL keys are often protected by a passphrase, but a developer may be understandably reluctant to specify that passphrase in a configuration file because it would be exposed in clear text. The developer would likely prefer to configure the IceSSL plug-in with a password callback instead; however, this must be done before the plug-in is initialized and attempts to load the SSL key. The solution is to configure the Ice run time so that it postpones the initialization of its plug-ins:

```
Ice.InitPlugins=0
```

When `Ice.InitPlugins` is set to zero, initializing plug-ins becomes the application's responsibility. The example below demonstrates how to perform this initialization:

C++

```
Ice::CommunicatorPtr ic = ...
Ice::PluginManagerPtr pm = ic->getPluginManager();
IceSSL::PluginPtr ssl = pm->getPlugin("IceSSL");
ssl->setPasswordPrompt(...);
pm->initializePlugins();
```

After obtaining the IceSSL plug-in and establishing the password callback, the application invokes `initializePlugins` on the plug-in manager object to commence plug-in initialization.

See Also

- IceSSL
- Ice.InitPlugins
- Ice.Plugin.*
- Ice.PluginLoadOrder

Copyright 2016, ZeroC, Inc.

Client-Side Features

This section presents all APIs and features that are specific to client applications.

Server-specific features are presented in Server-Side Features, while transverse features best understood while considering both sides of a client-server interaction are described in Client-Server Features.

Topics

- Proxies
- Request Contexts
- Invocation Timeouts
- Automatic Retries
- Oneway Invocations
- Datagram Invocations
- Batched Invocations

Copyright 2016, ZeroC, Inc.

Proxies

The introduction to proxies provided in Terminology describes a proxy as a local artifact that makes a remote invocation as easy to use as a regular function call. In fact, processing remote invocations is just one of a proxy's many responsibilities. A proxy also encapsulates the information necessary to contact the object, including its identity and addressing details such as endpoints. Proxy methods provide access to configuration and connection information, and act as factories for creating new proxies. Finally, a proxy initiates the establishment of a new connection when necessary.

Topics

- Obtaining Proxies
- Proxy Methods
- Proxy Endpoints
- Filtering Proxy Endpoints
- Proxy Defaults and Overrides
- Proxy-Based Load Balancing
- Proxy and Endpoint Syntax

See Also

- Terminology
- Object Identity
- Proxy Methods
- Connection Establishment

Copyright 2016, ZeroC, Inc.

Obtaining Proxies

This page describes the ways an application can obtain a proxy.

On this page:

- Obtaining a Proxy from a String
- Obtaining a Proxy from Properties
- Obtaining a Proxy using Factory Methods
- Obtaining a Proxy by Invoking Operations

Obtaining a Proxy from a String

The communicator operation `stringToProxy` creates a proxy from its stringified representation, as shown in the following C++ example:

C++

```
Ice::ObjectPrx p = communicator->stringToProxy("ident:tcp -p 5000");
```

Obtaining a Proxy from Properties

Rather than hard-coding a stringified proxy as the previous example demonstrated, an application can gain more flexibility by externalizing the proxy in a configuration property. For example, we can define a property that contains our stringified proxy as follows:

```
MyApp.Proxy=ident:tcp -p 5000
```

We can use the communicator operation `propertyToProxy` to convert the property's value into a proxy, as shown below in Java:

Java

```
Ice.ObjectPrx p = communicator.propertyToProxy("MyApp.Proxy");
```

As an added convenience, `propertyToProxy` allows you to define subordinate properties that configure the proxy's local settings. The properties below demonstrate this feature:

```
MyApp.Proxy=ident:tcp -p 5000
MyApp.Proxy.PreferSecure=1
MyApp.Proxy.EndpointSelection=Ordered
```

These additional properties simplify the task of customizing a proxy (as you can with proxy methods) without the need to change the application's code. The properties shown above are equivalent to the following statements:

Copyright 2016, ZeroC, Inc.

Java

```
Ice.ObjectPrx p = communicator.stringToProxy("ident:tcp -p 5000");
p = p.ice_preferSecure(true);
p = p.ice_endpointSelection(Ice.EndpointSelectionType.Ordered);
```

The list of supported proxy properties includes the most commonly-used proxy settings. The communicator prints a warning by default if it does not recognize a subordinate property. You can disable this warning using the property `Ice.Warn.UnknownProperties`.

Note that proxy properties can themselves have proxy properties. For example, the following sets the `PreferSecure` property on the default locator's router:

```
Ice.Default.Locator.Router.PreferSecure=1
```

Obtaining a Proxy using Factory Methods

Proxy factory methods allow you to modify aspects of an existing proxy. Since proxies are immutable, factory methods always return a new proxy if the desired modification differs from the proxy's current configuration. Consider the following C# example:

C#

```
Ice.ObjectPrx p = communicator.stringToProxy("...");
p = p.ice_oneway();
```

`ice_oneway` is considered a factory method because it returns a proxy configured to use oneway invocations. If the original proxy uses a different invocation mode, the return value of `ice_oneway` is a new proxy object.

The `checkedCast` and `uncheckedCast` methods can also be considered factory methods because they return new proxies that are narrowed to a particular Slice interface. A call to `checkedCast` or `uncheckedCast` typically follows the use of other factory methods, as shown below:

C#

```
Ice.ObjectPrx p = communicator.stringToProxy("...");
Ice.LocatorPrx loc = Ice.LocatorPrxHelper.checkedCast(p.ice_secure(true));
```

Note however that, once a proxy has been narrowed to a Slice interface, it is not normally necessary to perform another down-cast after using a factory method. For example, we can rewrite this example as follows:

Copyright 2016, ZeroC, Inc.

C#

```
Ice.ObjectPrx p = communicator.stringToProxy("...");
Ice.LocatorPrx loc = Ice.LocatorPrxHelper.checkedCast(p);
loc = (Ice.LocatorPrx)p.ice_secure(true);
```

A language-specific cast may be necessary, as shown here for C#, because the factory methods are declared to return the type `ObjectPrx`, but the proxy object itself retains its narrowed type. The only exceptions are the factory methods `ice_facet` and `ice_identity`. Calls to either of these methods may produce a proxy for an object of an unrelated type, therefore they return a base proxy that you must subsequently down-cast to an appropriate type.

Obtaining a Proxy by Invoking Operations

An application can also obtain a proxy as the result of an Ice invocation. Consider the following Slice definitions:

Slice

```
interface Account { ... };
interface Bank {
    Account* findAccount(string id);
};
```

Invoking the `findAccount` operation returns a proxy for an `Account` object. There is no need to use `checkedCast` or `uncheckedCast` on this proxy because it has already been narrowed to the `Account` interface. The C++ code below demonstrates how to invoke `findAccount`:

C++

```
BankPrx bank = ...
AccountPrx acct = bank->findAccount(id);
```

Of course, the application must have already obtained a proxy for the bank object using one of the techniques shown above.

See Also

- Communicators
- Proxy and Endpoint Syntax
- Proxy Methods
- Proxy Properties
- Ice.Warn.*

Copyright 2016, ZeroC, Inc.

Proxy Methods

Although the core proxy functionality is supplied by a language-specific base class, we can describe the proxy methods in terms of Slice operations as shown below:

Slice

```
bool ice_isA(string id);
void ice_ping();
StringSeq ice_ids();
string ice_id();
Communicator ice_getCommunicator();
string ice_toString();
Object* ice_identity(Identity id);
Identity ice_getIdentity();
Object* ice_adapterId(string id);
string ice_getAdapterId();
Object* ice_endpoints(EndpointSeq endpoints);
EndpointSeq ice_getEndpoints();
Object* ice_endpointSelection(EndpointSelectionType t);
EndpointSelectionType ice_getEndpointSelection();
Object* ice_context(Context ctx);
Context ice_getContext();
Object* ice_defaultContext();
Object* ice_facet(string facet);
string ice_getFacet();
Object* ice_twoway();
bool ice_isTwoway();
Object* ice_oneway();
bool ice_isOneway();
Object* ice_batchOneway();
bool ice_isBatchOneway();
Object* ice_datagram();
bool ice_isDatagram();
Object* ice_batchDatagram();
bool ice_isBatchDatagram();
Object* ice_secure(bool b);
bool ice_isSecure();
EncodingVersion ice_getEncodingVersion();
Object* ice_encodingVersion(EncodingVersion v);
Object* ice_preferSecure(bool b);
bool ice_isPreferSecure();
Object* ice_compress(bool b);
Object* ice_timeout(int timeout);
Object* ice_router(Router* rtr);
Router* ice_getRouter();
Object* ice_locator(Locator* loc);
Locator* ice_getLocator();
Object* ice_locatorCacheTimeout(int seconds);
int ice_getLocatorCacheTimeout();
Object* ice_collocationOptimized(bool b);
bool ice_isCollocationOptimized();
Object* ice_invocationTimeout(int timeout);
int ice_getInvocationTimeout();
```


Copyright 2016, ZeroC, Inc.

```
Object* ice_connectionId(string id);
string ice_getConnectionId();
Connection ice_getConnection();
Connection ice_getCachedConnection();
Object* ice_connectionCached(bool b);
bool ice_isConnectionCached();
void ice_flushBatchRequests();
```


Copyright 2016, ZeroC, Inc.

```
bool ice_invoke(string operation, OperationMode mode,
                ByteSeq inParams, out ByteSeq outParams);
```

These methods can be categorized as follows:

- Remote inspection: methods that return information about the remote object. These methods make remote invocations and therefore accept an optional trailing argument of type `Ice::Context`.
- Local inspection: methods that return information about the proxy's local configuration.
- Factory: methods that return new proxy instances configured with different features.
- Request processing: methods that flush batch requests and send "dynamic" Ice invocations.

Proxies are immutable, so factory methods allow an application to obtain a new proxy with the desired configuration. Factory methods essentially clone the original proxy and modify one or more features of the new proxy.

Many of the factory methods are not supported by fixed proxies, which are used in conjunction with bidirectional connections. Attempting to invoke one of these methods causes the Ice run time to raise `FixedProxyException`.

The core proxy methods are explained in greater detail in the following table:

Method	Description	Remote
`ice_isA`	Returns `true` if the remote object supports the type indicated by the `id` argument, otherwise `false`. This method can only be invoked on a twoway proxy.	Yes
`ice_ping`	Determines whether the remote object is reachable. Does not return a value.	Yes
`ice_ids`	Returns the type IDs of the types supported by the remote object. The return value is an array of strings. This method can only be invoked on a twoway proxy.	Yes
`ice_id`	Returns the type ID of the most-derived type supported by the remote object. This method can only be invoked on a twoway proxy.	Yes
`ice_getCommunicator`	Returns the communicator that was used to create this proxy.	No
`ice_toString`	Returns the string representation of the proxy.	No
`ice_identity`	Returns a new proxy having the given identity.	No
`ice_getIdentity`	Returns the identity of the Ice object represented by the proxy.	No
`ice_adapterId`	Returns a new proxy having the given adapter ID.	No
`ice_getAdapterId`	Returns the proxy's adapter ID, or an empty string if no adapter ID is configured.	No
`ice_endpoints`	Returns a new proxy having the given endpoints.	No
`ice_getEndpoints`	Returns a sequence of `Endpoint` objects representing the proxy's endpoints.	No
`ice_endpointSelection`	Returns a new proxy having the given endpoint selection policy (random or ordered).	No
`ice_getEndpointSelection`	Returns the endpoint selection policy for the proxy.	No
`ice_context`	Returns a new proxy having the given request context.	No
`ice_getContext`	Returns the request context associated with the proxy.	No
`ice_facet`	Returns a new proxy having the given facet name.	No
`ice_getFacet`	Returns the name of the facet associated with the proxy, or an empty string if no facet has been set.	No
`ice_twoway`	Returns a new proxy for making twoway invocations.	No
`ice_isTwoway`	Returns `true` if the proxy uses twoway invocations, otherwise `false`.	No
`ice_oneway`	Returns a new proxy for making oneway invocations.	No

Copyright 2016, ZeroC, Inc.

`ice_isOneway`	Returns `true` if the proxy uses oneway invocations, otherwise `false`.	No
`ice_batchOneway`	Returns a new proxy for making batch oneway invocations.	No
`ice_isBatchOneway`	Returns `true` if the proxy uses batch oneway invocations, otherwise `false`.	No
`ice_datagram`	Returns a new proxy for making datagram invocations.	No
`ice_isDatagram`	Returns `true` if the proxy uses datagram invocations, otherwise `false`.	No
`ice_batchDatagram`	Returns a new proxy for making batch datagram invocations.	No
`ice_isBatchDatagram`	Returns `true` if the proxy uses batch datagram invocations, otherwise `false`.	No
`ice_secure`	Returns a new proxy whose endpoints may be filtered depending on the boolean argument. If `true`, only endpoints using secure transports are allowed, otherwise all endpoints are allowed.	No
`ice_getEncodingVersion`	Returns the encoding version that is used to encode requests invoked on this proxy.	No
`ice_encodingVersion`	Returns a new proxy that uses the given encoding version to encode requests. Raises `UnsupportedEncodingException` if the version is invalid.	No
`ice_isSecure`	Returns `true` if the proxy uses only secure endpoints, otherwise `false`.	No
`ice_preferSecure`	Returns a new proxy whose endpoints are filtered depending on the boolean argument. If `true`, endpoints using secure transports are given precedence over endpoints using non-secure transports. If `false`, the default behavior gives precedence to endpoints using non-secure transports.	No
`ice_isPreferSecure`	Returns `true` if the proxy prefers secure endpoints, otherwise `false`.	No
`ice_compress`	Returns a new proxy whose protocol compression capability is determined by the boolean argument. If `true`, the proxy uses protocol compression if it is supported by the endpoint. If `false`, protocol compression is never used.	No
`ice_timeout`	Returns a new proxy whose endpoints all have the given connection timeout value in milliseconds. A value of `-1` disables timeouts.	No
`ice_router`	Returns a new proxy configured with the given router proxy.	No
`ice_getRouter`	Returns the router that is configured for the proxy (null if no router is configured).	No
`ice_locator`	Returns a new proxy with the specified locator.	No
`ice_getLocator`	Returns the locator that is configured for the proxy (null if no locator is configured).	No
`ice_locatorCacheTimeout`	Returns a new proxy with the specified locator cache timeout in seconds. When binding a proxy to an endpoint, the run time caches the proxy returned by the locator and uses the cached proxy while the cached proxy has been in the cache for less than the timeout. Proxies older than the timeout cause the run time to rebind via the locator. A value of 0 disables caching entirely, and a value of -1 means that cached proxies never expire. The default value is -1.	No
`ice_getLocatorCacheTimeout`	Returns the locator cache timeout value in seconds.	No
`ice_collocationOptimized`	Returns a new proxy configured for collocation optimization. If `true`, collocated optimizations are enabled. The default value is `true`.	No
`ice_isCollocationOptimized`	Returns `true` if the proxy uses collocation optimization, otherwise `false`.	No
`ice_invocationTimeout`	Returns a new proxy configured with the specified invocation timeout in milliseconds. A value of `-1` means an invocation never times out. A value of `-2` provides backward compatibility with Ice versions prior to 3.6 by disabling invocation timeouts and using connection timeouts to wait for the response of an invocation.	No
`ice_getInvocationTimeout`	Returns the invocation timeout value in milliseconds.	No

Copyright 2016, ZeroC, Inc.

`ice_connectionId`	Returns a new proxy having the given connection ID.	No
`ice_getConnectionId`	Returns the connection ID, or an empty string if no connection ID has been configured.	No
`ice_getConnection` `begin_ice_getConnection`	Returns an object representing the connection used by the proxy. If the proxy is not currently associated with a connection, the Ice run time attempts to establish a connection first, which means this method can potentially block while network operations are in progress. Use the asynchronous method to avoid blocking.	No
`ice_getCachedConnection`	Returns an object representing the connection used by the proxy, or null if the proxy is not currently associated with a connection.	No
`ice_connectionCached`	Enables or disables connection caching for the proxy.	No
`ice_isConnectionCached`	Returns `true` if the proxy uses connection caching, otherwise `false`.	No
`ice_flushBatchRequests` `begin_ice_flushBatchRequests`	Sends a batch of operation invocations synchronously or asynchronously.	Yes
`ice_invoke` `begin_ice_invoke`	Allows dynamic invocation of an operation without the need for compiled Slice definitions. Requests can be sent synchronously or asynchronously.	Yes

See Also

- Request Contexts
- Oneway Invocations
- Batched Invocations
- Versioning
- Connection Timeouts
- Connection Establishment
- Using Connections
- Dynamic Invocation and Dispatch
- Asynchronous Dynamic Invocation and Dispatch
- Bidirectional Connections
- Glacier2
- Data Encoding

Copyright 2016, ZeroC, Inc.

Proxy Endpoints

Proxy endpoints are the client-side equivalent of object adapter endpoints. A proxy endpoint identifies the protocol information used to contact a remote object, as shown in the following example:

```
tcp -h frosty.zeroc.com -p 10000
```

This endpoint states that an object is reachable via TCP on the host `frosty.zeroc.com` and the port `10000`.

A proxy must have, or be able to obtain, at least one endpoint in order to be useful. A *direct proxy* contains one or more endpoints:

```
MyObject:tcp -h frosty.zeroc.com -p 10000:ssl -h frosty.zeroc.com -p 10001
```

In this example the object with the identity `MyObject` is available at two separate endpoints, one using TCP and the other using SSL.

If a direct proxy does not contain the `-h` option (that is, no host is specified), the Ice run time uses the value of the `Ice.Default.Host` property. If `Ice.Default.Host` is not defined, the localhost interface is used.

An *indirect proxy* uses a locator to retrieve the endpoints dynamically. One style of indirect proxy contains an adapter identifier:

```
MyObject @ MyAdapter
```

When this proxy requires the endpoints associated with `MyAdapter`, it requests them from the locator.

See Also

- Terminology
- Object Adapter Endpoints
- Locators

Copyright 2016, ZeroC, Inc.

Filtering Proxy Endpoints

A proxy's configuration determines how its endpoints are used. For example, a proxy configured for secure communication will only use endpoints having a secure protocol, such as SSL.

The factory methods listed in the table below allow applications to manipulate endpoints indirectly. Calling one of these methods returns a new proxy whose endpoints are used in accordance with the proxy's new configuration.

Method	Description
`ice_secure`	Selects only endpoints using a secure protocol (e.g., SSL).
`ice_datagram`	Selects only endpoints using a datagram protocol (e.g., UDP).
`ice_batchDatagram`	Selects only endpoints using a datagram protocol (e.g., UDP).
`ice_twoway`	Selects only endpoints capable of making twoway invocations (e.g., TCP, SSL). For example, this disables datagram endpoints.
`ice_oneway`	Selects only endpoints capable of making reliable oneway invocations (e.g., TCP, SSL). For example, this disables datagram endpoints.
`ice_batchOneway`	Selects only endpoints capable of making reliable oneway batch invocations (e.g., TCP, SSL). For example, this disables datagram endpoints.

Upon return, the set of endpoints in the new proxy is unchanged from the old one. However, the new proxy's configuration drives a filtering process that the Ice run time performs during connection establishment.

The factory methods do not raise an exception if they produce a proxy with no viable endpoints. For example, the C++ statement below creates such a proxy:

C++

```
proxy = comm->stringToProxy("id:tcp -p 10000")->ice_datagram();
```

It is always possible that a proxy could become viable after additional factory methods are invoked, therefore the Ice run time does not raise an exception until connection establishment is attempted. At that point, the application can expect to receive `NoEndpointException` if the filtering process eliminates all endpoints.

An application can also create a proxy with a specific set of endpoints using the `ice_endpoints` factory method, whose only argument is a sequence of `Ice::Endpoint` objects. At present, an application is not able to create new instances of `Ice::Endpoint`, but rather can only incorporate instances obtained by calling `ice_getEndpoints` on a proxy. Note that `ice_getEndpoints` may return an empty sequence if the proxy has no endpoints, as is the case with an indirect proxy.

See Also

- Proxy Endpoints
- Proxy Methods
- Connection Establishment

Copyright 2016, ZeroC, Inc.

Proxy Defaults and Overrides

It is important to understand how proxies are influenced by Ice configuration properties and settings. The relevant properties can be classified into two categories: defaults and overrides.

On this page:

- Proxy Default Properties
- Proxy Override Properties

Proxy Default Properties

Default properties affect proxies created as the result of an Ice invocation, or by calling `stringToProxy` or `propertyToProxy` on a communicator. These properties do not influence proxies created by proxy factory methods.

For example, suppose we define the following default property:

```
Ice.Default.PreferSecure=1
```

We can verify that the property has the desired affect using the following C++ code:

C++

```
Ice::ObjectPrx p = communicator->stringToProxy(...);
assert(p->ice_isPreferSecure());
```

Furthermore, we can verify that the property does not affect proxies returned by factory methods:

```
Ice::ObjectPrx p2 = p->ice_preferSecure(false);
assert(!p2->ice_isPreferSecure());
Ice::ObjectPrx p3 = p2->ice_oneway();
assert(!p3->ice_isPreferSecure());
```

Proxy Override Properties

Defining an override property causes the Ice run time to ignore any equivalent proxy setting and use the override property value instead. For example, consider the following property definition:

```
Ice.Override.Secure=1
```

This property instructs the Ice run time to use only secure endpoints, producing the same semantics as calling `ice_secure(true)` on every proxy. However, the property does not alter the settings of an existing proxy, but rather directs the Ice run time to use secure endpoints regardless of the proxy's security setting. We can verify that this is the case using the following C++ code:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx p = communicator->stringToProxy(...);
p = p->ice_secure(false);
assert(!p->ice_isSecure()); // The security setting is retained.
```

See Also

- Communicators
- Ice.Default.*
- Ice.Override.*

Copyright 2016, ZeroC, Inc.

Proxy-Based Load Balancing

Ice supports two types of load balancing:

- Locator-based load balancing
- Proxy-based load balancing

An application can use only one type or combine both to achieve the desired behavior. This page discusses proxy-based load balancing; refer to Locator Semantics for Clients for information on load balancing with a locator.

On this page:

- Proxy Connection Caching
- Proxies with Multiple Endpoints
- Per-Request Load Balancing

Proxy Connection Caching

Before we discuss load balancing, it's important to understand the relationship between proxies and connections.

By default a proxy remembers its connection and uses it for all invocations until that connection is closed. You can prevent a proxy from caching its connection by calling the `ice_connectionCached` proxy method with an argument of false. Once connection caching is disabled, each invocation on a proxy causes the Ice run time to execute its connection establishment process.

Note that each invocation on such a proxy does not necessarily cause the Ice run time to establish a new connection. It only means that the Ice run time does not assume that it can reuse the connection of the proxy's previous invocation. Whether the Ice run time actually needs to establish a new connection for the next invocation depends on several factors.

As with any feature, you should only use it when the benefits outweigh the risks. With respect to a proxy's connection caching behavior, there is certainly a small amount of computational overhead associated with executing the connection establishment process for each invocation, as well as the risk of significant overhead each time a new connection is actually created.

Proxies with Multiple Endpoints

A proxy can contain zero or more endpoints. A proxy with no endpoints typically means it's an indirect proxy that requires a location service t o convert the proxy's symbolic information into endpoints at run time. Otherwise, a direct proxy contains at least one endpoint.

Regardless of whether endpoints are specified directly or indirectly, a proxy having multiple endpoints implies at a minimum that the target object is available via multiple network interfaces, but often also means the object is replicated to improve scalability and reliability. Such a proxy provides a client with several load balancing options even when no location service is involved. The proxy's own configuration drives the run-time behavior, depending on how the client configures the proxy's endpoint selection type and connection caching settings.

For example, suppose that a proxy contains several endpoints. In its default configuration, a proxy uses the `Random` endpoint selection type and caches its connection. Upon the first invocation, the Ice run time selects one of the proxy's endpoints at random and uses that connection for all subsequent invocations until the connection is closed. For some applications, this form of load balancing may be sufficient.

Suppose now that we use the `Ordered` endpoint selection type instead. In this case, the Ice run time always attempts to establish connections using the endpoints in the order they appear in the proxy. Normally an application uses this configuration when there is a preferred order to the servers. Again, once connected, the Ice run time uses whichever connection was chosen indefinitely.

Per-Request Load Balancing

When we disable the connection caching behavior of a proxy with multiple endpoints, its semantics undergo a significant change. Using the `Random` endpoint selection type, the Ice run time selects one of the endpoints at random and establishes a connection to it if one is not already established, and this process is repeated *prior to each subsequent invocation*. This is called *per-request load balancing* because each request can potentially be directed to a different server.

Using the `Ordered` endpoint selection type is not as common in this scenario; its main purpose would be to fall back on a secondary server if the primary server is not available, but it causes the Ice run time to attempt to contact the primary server during each request.

Here's some code that shows how to configure the proxy:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx proxy = communicator->stringToProxy("hello:tcp -h 10.0.0.1 -p
2000:tcp -h 10.0.0.2 -p 2001");
proxy = proxy->ice_connectionCached(false);
proxy = proxy->ice_endpointSelection(Ice::Random);
// If also using a locator:
proxy = proxy->ice_locatorCacheTimeout(...);
```

We create a proxy with two endpoints, then use the factory methods to disable connection caching and set the endpoint selection type. The `Random` setting means that, after we've made a few invocations with the proxy, the Ice run time will have established connections to both endpoints. After incurring the initial expense of opening the connections, each subsequent invocation will randomly use one of the existing connections.

If you're also using a location service, you may want to modify the proxy's locator cache timeout to force the Ice run time in the client to query the locator more frequently.

See Also

- Terminology
- Proxy Methods
- Connection Establishment
- Active Connection Management
- Ice.Default.*
- Locator Semantics for Clients
- Load Balancing

Copyright 2016, ZeroC, Inc.

Proxy and Endpoint Syntax

On this page:

- Syntax for Stringified Proxies
- Syntax for Stringified Endpoints
  - Address Syntax
  - TCP Endpoint Syntax
  - UDP Endpoint Syntax
  - SSL Endpoint Syntax
  - WS Endpoint Syntax
  - WSS Endpoint Syntax
  - Opaque Endpoint Syntax

Syntax for Stringified Proxies

Synopsis

identity -f *facet* -e *encoding* -p *protocol* -t -o -O -d -D -s @ *adapter_id* : *endpoints*

Description

A stringified proxy consists of an identity, proxy options, and an optional object adapter identifier or endpoint list. White space (the space, tab (`\t`), line feed (`\n`), and carriage return (`\r`) characters) act as token delimiters; if a white space character appears as part of a component of a stringified proxy (such as the identity), it must be quoted or escaped as described below.

A proxy containing an identity with no endpoints, or an identity with an object adapter identifier, represents an indirect proxy that will be resolved using the Ice locator.

Proxy options configure the invocation mode:

`-f` *facet*	Select a facet of the Ice object.
`-e` *encoding*	Specify the Ice encoding version to use when an invocation on this proxy marshals parameters.
`-p` *protocol*	Specify the Ice protocol version to use when sending a request with this proxy.
`-t`	Configures the proxy for twoway invocations (default).
`-o`	Configures the proxy for oneway invocations.
`-O`	Configures the proxy for batch oneway invocations.
`-d`	Configures the proxy for datagram invocations.
`-D`	Configures the proxy for batch datagram invocations.
`-s`	Configures the proxy for secure invocations.

The proxy options `-t`, `-o`, `-O`, `-d`, and `-D` are mutually exclusive.

The object identity `identity` is structured as `[`*category*`/]`*name*, where the *category* component and slash separator are optional. If *identity* contains white space or either of the characters `:` or `@`, it must be enclosed in single or double quotes. The *category* and *name* components are UTF-8 strings that use the encoding described below. Any occurrence of a slash (`/`) in *category* or *name* must be escaped with a backslash (i.e., `\/`).

The *facet* argument of the `-f` option represents a facet name. If *facet* contains white space, it must be enclosed in single or double quotes. A facet name is a UTF-8 string that uses the encoding described below.

The object adapter identifier *adapter_id* is a UTF-8 string that uses the encoding described below. If *adapter_id* contains white space, it must be enclosed in single or double quotes.

Single or double quotes can be used to prevent white space characters from being interpreted as delimiters. Double quotes prevent interpretation of a single quote as an opening or closing quote, for example:

Copyright 2016, ZeroC, Inc.

```
"a string with a ' quote"
```

Single quotes prevent interpretation of a double quote as an opening or closing quote. For example:

```
'a string with a " quote'
```

Escape sequences such as `\b` are interpreted within single and double quotes.

UTF-8 strings are encoded using ASCII characters for the ordinal range 32--126 (inclusive). Characters outside this range must be encoded using escape sequences (`\b`, `\f`, `\n`, `\r`, `\t`) or octal notation (e.g., `\007`). Single and double quotes can be escaped using a backslash, as can the backslash itself.

If *`endpoints`* are specified, they must be separated with a colon (`:`) and formatted as described in the endpoint syntax. The order of endpoints in the stringified proxy is not necessarily the order in which connections are attempted during binding: when a stringified proxy is converted into a proxy instance, by default, the endpoint list is randomized as a form of load balancing. You can change this default behavior using the properties `Ice.Default.EndpointSelection` and *`name`*`.EndpointSelection`.

If the `-s` option is specified, only those endpoints that support secure invocations are considered during binding. If no valid endpoints are found, the application receives `Ice::NoEndpointException`.

Otherwise, if the `-s` option is not specified, the endpoint list is ordered so that non-secure endpoints have priority over secure endpoints during binding. In other words, connections are attempted on all non-secure endpoints before any secure endpoints are attempted.

The `-e` and `-p` options specify the encoding and protocol versions supported by the target object, respectively. The Ice run time in the client is responsible for using a compatible version of the encoding or protocol, where *compatible* means the same major version and a minor version that is equal to or less than the specified minor version. If the `-e` option is not defined, the proxy uses the default version specified by `Ice.Default.EncodingVersion`.

If an unknown option is specified, or the stringified proxy is malformed, the application receives `Ice::ProxyParseException`. If an endpoint is malformed, the application receives `Ice::EndpointParseException`.

Syntax for Stringified Endpoints

Synopsis

`endpoint` `:` *`endpoint`*

Description

An endpoint list comprises one or more endpoints separated by a colon (`:`).

An endpoint has the following format:

`protocol option`

The supported protocols are `tcp`, `udp`, `ssl`, `ws`, `wss`, and `default`. If `default` is used, it is replaced by the value of the `Ice.Default.Protocol` property. If an endpoint is malformed, or an unknown protocol is specified, the application receives `Ice::EndpointParseException`. The `ssl` and `wss` protocols are only available if the IceSSL plug-in is installed.

Ice uses endpoints for two similar but distinct purposes:

1. In a client context (that is, in a proxy), endpoints determine how Ice establishes a connection to a server.
2. In a server context (that is, in an object adapter's configuration), endpoints define the addresses and transports over which new incoming connections are accepted. These endpoints are also embedded in the proxies created by the object adapter, unless a separate set of "published" endpoints are explicitly configured.

The sections that follow discuss the addressing component of endpoints, as well as the protocols and their supported options.

See Object Adapter Endpoints for examples.

Copyright 2016, ZeroC, Inc.

Address Syntax

Synopsis

```
host : hostname | x.x.x.x (IPv4)
host : hostname | ":x:x:x:x:x:x:x" (IPv6)
```

Description

Ice supports Internet Protocol (IP) versions 4 and 6 in all language mappings.

Support for these protocols is configured using the properties `Ice.IPv4` and `Ice.IPv6` (both enabled by default).

In the endpoint descriptions below, the *host* parameter represents either a host name that is resolved via the Domain Name System (DNS), an IPv4 address in dotted quad notation, or an IPv6 address in 128-bit hexadecimal format and enclosed in double quotes. Due to limitation of the DNS infrastructure, host and domain names are restricted to the ASCII character set.

The presence (or absence) of the *host* parameter has a significant influence on the behavior of the Ice run time. The table below describes these semantics:

Value	Client Semantics	Server Semantics
None	If *host* is not specified in a proxy, Ice uses the value of the `Ice.Default.Host` property. If that property is not defined, outgoing connections are only attempted over loopback interfaces.	If *host* is not specified in an object adapter endpoint, Ice uses the value of the `Ice.Default.Host` property. If that property is not defined, the adapter behaves as if the wildcard symbol * was specified (see below).
Host name	The host name is resolved via DNS. Outgoing connections are attempted to each address returned by the DNS query.	The host name is resolved via DNS, and the object adapter listens on the network interfaces corresponding to each address returned by the DNS query. The specified host name is embedded in proxies created by the adapter.
IPv4 address	An outgoing connection is attempted to the given address.	The object adapter listens on the network interface corresponding to the address. The specified address is embedded in proxies created by the adapter.
IPv6 address	An outgoing connection is attempted to the given address.	The object adapter listens on the network interface corresponding to the address. The specified address is embedded in proxies created by the adapter.
`0.0.0.0` (IPv4)	A "wildcard" IPv4 address that causes Ice to try all local interfaces when establishing an outgoing connection.	The adapter listens on all IPv4 network interfaces (including the loopback interface), that is, binds to INADDR_ANY for IPv4. Endpoints for all addresses except loopback are published in proxies (unless loopback is the only available interface, in which case only loopback is published).
`"::"` (IPv6)	A "wildcard" IPv6 address that causes Ice to try all local interfaces when establishing an outgoing connection.	Equivalent to * (see below).
* (IPv4, IPv6)	Not supported in proxies.	The adapter listens on all network interfaces (including the loopback interface), that is, binds to INADDR_ANY for the enabled protocols (IPv4 and/or IPv6). Endpoints for all addresses except loopback and IPv6 link-local are published in proxies (unless loopback is the only available interface, in which case only loopback is published).

There is one additional benefit in specifying a wildcard address for *host* (or not specifying it at all) in an object adapter's endpoint: if the list of network interfaces on a host may change while the application is running, using a wildcard address for *host* ensures that the object adapter automatically includes the updated interfaces. Note however that the list of published endpoints is not changed automatically; rather, the application must explicitly refresh the object adapter's endpoints. For diagnostic purposes, you can set the configuration property `Ice.Trace.Network=3` to cause Ice to log the current list of local addresses that it is substituting for the wildcard address.

When IPv4 and IPv6 are enabled, an object adapter endpoint that uses an IPv6 (or * wildcard) address can accept both IPv4 and IPv6 connections.

Java's default network stack always accepts both IPv4 and IPv6 connections regardless of the settings of `Ice.IPv6`. Thus, in Java, an object adapter endpoint that uses the IPv4 wildcard will accept both IPv4 and IPv6 connections. You can configure the Java run time to use only IPv4 by starting your application with the following JVM option:

Copyright 2016, ZeroC, Inc.

Java

```
java -Djava.net.preferIPv4Stack=true ...
```

TCP Endpoint Syntax

Synopsis

```
tcp -h host -p port -t timeout -z --sourceAddress addr
```

Description

A `tcp` endpoint supports the following options:

Option	Description	Client Semantics	Server Semantics
`-h` *`host`*	Specifies the host name or IP address of the endpoint. If not specified, the value of `Ice.Default.Host` is used instead.	See Address Syntax.	See Address Syntax.
`-p` *`port`*	Specifies the port number of the endpoint.	Determines the port to which a connection attempt is made (required).	The port will be selected by the operating system if this option is not specified or *`port`* is zero.
`-t` *`timeout`*	Specifies the endpoint timeout in milliseconds.	The value for *`timeout`* must either be `infinite` to specify no timeout, or an integer greater than zero representing the timeout in milliseconds used by the client to open or close connections and to read or write data. If a timeout occurs, the application receives `Ice::TimeoutException`. If this option is not specified, it defaults to the value of `Ice.Default.Timeout`.	The value for *`timeout`* must either be `infinite` to specify no timeout, or an integer greater than zero representing the timeout in milliseconds used by the server to accept or close connections and to read or write data (see Timeouts in Object Adapter Endpoints and Connection Timeouts). *`timeout`* also controls the timeout that is published in proxies created by the object adapter. If this option is not specified, it defaults to the value of `Ice.Default.Timeout`.
`-z`	Specifies bzip2 compression.	Determines whether compressed requests are sent.	Determines whether compression is advertised in proxies created by the adapter.
`--sourceAddress` *`addr`*	Binds outgoing socket connections to the network interface associated with *`addr`*.	The value for *`addr`* must be a numeric IPv4 or IPv6 address. If this option is not specified, it defaults to the value of `Ice.Default.SourceAddress`. This feature is not supported on WinRT.	Not supported

UDP Endpoint Syntax

Synopsis

```
udp -h host -p port -z --ttl TTL --interface INTF --sourceAddress addr
```

Description

A `udp` endpoint supports either unicast or multicast delivery; the address resolved by *`host`* argument determines the delivery mode. To use multicast in IPv4, select an IP address in the range `233.0.0.0 to 239.255.255.255`. In IPv6, use an address that begins with `ff`, such as `ff01::1:1`.

Copyright 2016, ZeroC, Inc.

A udp endpoint supports the following options:

Option	Description	Client Semantics	Server Semantics
-h *host*	Specifies the host name or IP address of the endpoint. If not specified, the value of Ice.Default.Host is used instead.	See Address Syntax.	See Address Syntax.
-p *port*	Specifies the port number of the endpoint.	Determines the port to which datagrams are sent (required).	The port will be selected by the operating system if this option is not specified or port is zero.
-z	Specifies bzip2 compression.	Determines whether compressed requests are sent.	Determines whether compression is advertised in proxies created by the adapter.
--ttl *TTL*	Specifies the time-to-live (also known as "hops") of multicast messages.	Determines whether multicast messages are forwarded beyond the local network. If not specified, or the value of *TTL* is -1, multicast messages are not forwarded. The maximum value is 255.	N/A
--interface *INTF*	Specifies the network interface or group for multicast messages (see below).	Selects the network interface for outgoing multicast messages. If not specified, multicast messages are sent using the default interface.	Selects the network interface to use when joining the multicast group. If not specified, the group is joined on the default network interface.
--sourceAddress *addr*	Binds outgoing socket connections to the network interface associated with *addr*.	The value for *addr* must be a numeric IPv4 or IPv6 address. If this option is not specified, it defaults to the value of Ice.Default.SourceAddress. This feature is not supported on WinRT.	Not supported

Deprecated options
With the 1.0 encoding, UDP endpoints supported 2 additional options: the -e *major.minor* and -v *major.minor* options. These 2 options specified which encoding and protocol was supported by the endpoint. These two options are deprecated with the 1.1 encoding and are ignored (a deprecation warning will be emitted when parsed by the communicator stringToProxy method). The supported protocol and encoding is specified on the proxy with the 1.1 encoding.

Multicast Interfaces

When *host* denotes a multicast address, the --interface *INTF* option selects a particular network interface to be used for communication. The format of *INTF* depends on the language and IP version:

- C++ and .NET
  INTF can be an interface name, such as eth0, or an IP address. If using IPv6, it can also be an interface index. Interface names on Windows may contain spaces, such as Local Area Connection, therefore they must be enclosed in double quotes.
- Java
  INTF can be an interface name, such as eth0, or an IP address. On Windows, Java maps interface names to Unix-style nicknames.

SSL Endpoint Syntax

Synopsis

```
ssl -h host -p port -t timeout -z --sourceAddress addr
```

Description

An ssl endpoint supports the same options as for tcp endpoints.

WS Endpoint Syntax

Copyright 2016, ZeroC, Inc.

Synopsis

```
ws -r resource -h host -p port -t timeout -z --sourceAddress addr
```

Description

A `ws` (WebSocket) endpoint supports all `tcp` endpoint options in addition to the following:

Option	Description	Client Semantics	Server Semantics
`-r resource`	A URI specifying the resource associated with this endpoint. If not specified, the default value is `/`.	The value for *`resource`* is passed as the target for `GET` in the WebSocket upgrade request.	The web server configuration must direct the given *`resource`* to this endpoint.

WSS Endpoint Syntax

Synopsis

```
wss -r resource -h host -p port -t timeout -z --sourceAddress addr
```

Description

A `wss` (Secure WebSocket) endpoint supports all `ssl` endpoint options in addition to the following:

Option	Description	Client Semantics	Server Semantics
`-r resource`	A URI specifying the resource associated with this endpoint. If not specified, the default value is `/`.	The value for *`resource`* is passed as the target for `GET` in the WebSocket upgrade request.	The web server configuration must direct the given *`resource`* to this endpoint.

Opaque Endpoint Syntax

Synopsis

```
opaque -t type -e encoding -v value
```

Description

Proxies can contain endpoints that are not universally understood by Ice processes. For example, a proxy can contain an SSL endpoint; if that proxy is marshaled to a receiver without the IceSSL plug-in, the SSL endpoint does not make sense to the receiver.

Ice preserves such unknown endpoints when they are received over the wire. For the preceding example, if the receiver remarshals the proxy and sends it back to an Ice process that does have the IceSSL plug-in, that process can invoke on the proxy using its SSL transport. This mechanism allows proxies containing endpoints for arbitrary transports to pass through processes that do not understand these endpoints without losing information.

If an Ice process stringifies a proxy containing an unknown endpoint, it writes the endpoint as an opaque endpoint. For example:

```
opaque -t 2 -e 1.0 -v CTEyNy4wLjAuMREnAAD/////AA==
```

This is how a process without the IceSSL plug-in stringifies an SSL endpoint. When a process with the IceSSL plug-in unstringifies this endpoint and converts it back into a string, it produces:

```
ssl -h 127.0.0.1 -p 10001
```

An `opaque` endpoint supports the following options:

Copyright 2016, ZeroC, Inc.

Option	Description
`-t` *`type`*	Specifies the transport for the endpoint. Transports are indicated by positive integers (1 for TCP, 2 for SSL, and 3 for UDP).
`-e` *`encoding`*	Specifies the encoding version used to encode the endpoint marshaled data.
`-v` *`value`*	Specifies the marshaled encoding of the endpoint in base-64 encoding.

Exactly one each of the `-t` and `-v` options must be present in an opaque endpoint. If `-e` is not specified, the current encoding supported by the Ice runtime is assumed.

See Also

- The Ice Protocol
- Object Adapter Endpoints
- Connection Timeouts
- Ice.Default.*

Copyright 2016, ZeroC, Inc.

Request Contexts

Methods on a proxy are overloaded with a trailing parameter representing a *request context*. The Slice definition of this parameter is as follows:

Slice

```
module Ice {
   dictionary<string, string> Context;
};
```

As you can see, a context is a dictionary that maps strings to strings or, conceptually, a context is a collection of name-value pairs. The contents of this dictionary (if any) are implicitly marshaled with every request to the server, that is, if the client populates a context with a number of name-value pairs and uses that context for an invocation, the name-value pairs that are sent by the client are available to the server.

On the server side, the operation implementation can access the received `Context` via the `ctx` member of the `Ice::Current` parameter and extract the name-value pairs that were sent by the client.

Context names beginning with an underscore are reserved for use by Ice.

Topics

- Explicit Request Contexts
- Per-Proxy Request Contexts
- Implicit Request Contexts
- Design Considerations for Request Contexts

See Also

- The Current Object

Copyright 2016, ZeroC, Inc.

Explicit Request Contexts

Request contexts provide a means of sending an unlimited number of parameters from client to server without having to mention these parameters in the signature of an operation. For example, consider the following definition:

Slice

```
struct Address {
    // ...
};

interface Person {
    string setAddress(Address a);
    // ...
};
```

Assuming that the client has a proxy to a `Person` object, it could do something along the following lines:

C++

```
PersonPrx p = ...;
Address a = ...;

Ice::Context ctx;
ctx["write policy"] = "immediate";

p->setAddress(a, ctx);
```

In Java, the same code would looks as follows:

Java

```
PersonPrx p = ...;
Address a = ...;

java.util.Map<String, String> ctx = new java.util.HashMap<String, String>();
ctx.put("write policy", "immediate");

p.setAddress(a, ctx);
```

In C#, the code is almost identical:

Copyright 2016, ZeroC, Inc.

C#

```
using System.Collections.Generic;

PersonPrx p = ...;
Address a = ...;

Dictionary<string, string> ctx = new Dictionary<string, string>();
ctx["write policy"] = "immediate";

p.setAddress(a, ctx);
```

On the server side, we can extract the policy value set from the `Current` object to influence how the implementation of `setAddress` works. A C++ implementation might look like this:

C++

```
void
PersonI::setAddress(const Address& a, const Ice::Current& c)
{
    Ice::Context::const_iterator i = c.ctx.find("write policy");
    if (i != c.ctx.end() && i->second == "immediate") {

        // Update the address details and write through to the
        // data base immediately...

    } else {

        // Write policy was not set (or had a bad value), use
        // some other database write strategy.
    }
}
```

For this example, the server examines the value of the context with the key `"write policy"` and, if that value is `"immediate"`, writes the update sent by the client straight away; if the write policy is not set or contains a value that is not recognized, the server presumably applies a more lenient write policy (such as caching the update in memory and writing it later).

See Also

- Request Contexts
- The Current Object

Copyright 2016, ZeroC, Inc.

Per-Proxy Request Contexts

Instead of passing a context explicitly with an invocation, you can also use a *per-proxy context*. Per-proxy contexts allow you to set a context on a particular proxy once and, thereafter, whenever you use that proxy to invoke an operation, the previously-set context is sent with each invocation.

On this page:

- Configuring a Per-Proxy Request Context Programmatically
- Configuring a Per-Proxy Request Context using Properties

Configuring a Per-Proxy Request Context Programmatically

The proxy methods `ice_context` and `ice_getContext` set and retrieve the context, respectively. The Slice definitions of these methods would look as follows:

Slice

```
Object* ice_context(Context ctx);
Context ice_getContext();
```

`ice_context` creates a new proxy that stores the given context. Calling `ice_getContext` returns the stored context, or an empty dictionary if no per-proxy context has been configured for the proxy.

Here is an example in C++:

C++

```
Ice::Context ctx;
ctx["write policy"] = "immediate";

PersonPrx p1 = ...;
PersonPrx p2 = p1->ice_context(ctx);

Address a = ...;

p1->setAddress(a);       // Sends no context

p2->setAddress(a);       // Sends ctx implicitly

Ice::Context ctx2;
ctx2["write policy"] = "delayed";

p2->setAddress(a, ctx2); // Sends ctx2 instead
```

As the example illustrates, once we have created the `p2` proxy, any invocation via `p2` automatically sends the configured context. The final line of the example shows that it is also possible to explicitly send a context for an invocation even if the proxy is configured with a context — an explicit context always overrides any per-proxy context.

Note that, once you have set a per-proxy context, that context becomes immutable: if you subsequently change the context you have passed to `ice_context`, such a change does not affect the per-proxy context of any proxies you previously created with that context because each proxy on which you set a per-proxy context stores its own copy of the dictionary.

Configuring a Per-Proxy Request Context using Properties

Copyright 2016, ZeroC, Inc.

As of Ice 3.6, you can also configure a context with proxy properties when you use the communicator operation `propertyToProxy`. Suppose we obtain a proxy like this:

C++

```
Ice::CommunicatorPtr communicator = ...;
Ice::ObjectPrx proxy = communicator->propertyToProxy("PersonProxy");
```

We can statically configure a context for this proxy using the following properties:

```
PersonProxy=person:tcp -p 5000
PersonProxy.Context.write policy=immediate
PersonProxy.Context.failure mode=persistent
```

The `Context` property has the form *`name`*`.Context.`*`key`*`=`*`value`*, where *`key`* and *`value`* can be any legal property symbols.

The proxy returned by `propertyToProxy` already contains the context key/value pairs specified in the configuration properties. To make any modifications to the context at run time, you'll need to retrieve the proxy's context dictionary using `ice_getContext`, modify the dictionary as necessary, and finally obtain a new proxy by calling `ice_context`, as we described above.

See Also

- Explicit Request Contexts
- Proxy Methods
- Proxy Properties

Copyright 2016, ZeroC, Inc.

Implicit Request Contexts

On this page:

- Using Implicit Request Contexts
- Scope of the Implicit Context

Using Implicit Request Contexts

In addition to explicit and per-proxy request contexts, you can also establish an implicit context on a communicator. This implicit context is sent with all invocations made via proxies created by that communicator, provided that you do not supply an explicit context with the call.

Access to this implicit context is provided by the `Communicator` interface:

Slice

```
module Ice {
    local interface Communicator
    {
        ImplicitContext getImplicitContext();

        // ...
    };
};
```

`getImplicitContext` returns the implicit context object. If a communicator has no implicit context, the operation returns null.

You can manipulate the contents of the implicit context via the `ImplicitContext` interface:

Slice

```
local interface ImplicitContext
{
    Context getContext();
    void    setContext(Context newContext);

    string get(string key);
    string put(string key, string value);
    string remove(string key);
    bool   containsKey(string key);
};
```

The `getContext` operation returns the currently-set context dictionary. The `setContext` operation replaces the currently-set context in its entirety.

The remaining operations allow you to manipulate specific entries:

- `get`
  This operation returns the value associated with `key`. If `key` was not previously set, the operation returns the empty string.
- `put`
  This operation adds the key-value pair specified by `key` and `value`. It returns the previous value associated with `key`; if no value was previously associated with `key`, it returns the empty string. It is legal to add the empty string as a value.
- `remove`
  This operation removes the key-value pair specified by `key`. It returns the previously-set value (or the empty string if `key` was not

Copyright 2016, ZeroC, Inc.

previously set).

- `containsKey`
  This operation returns true if `key` is currently set and false, otherwise. You can use this operation to distinguish between a key-value pair that was explicitly added with an empty string as a value, and a key-value pair that was never added at all.

Scope of the Implicit Context

You establish the implicit context on a communicator by setting a property, `Ice.ImplicitContext`. This property controls whether a communicator has an implicit context and, if so, at what scope the context applies. The property can be set to the following values:

- `None`
  With this setting (or if `Ice.ImplicitContext` is not set at all), the communicator has no implicit context, and `getImplicitContext` returns null.
- `Shared`
  The communicator has a single implicit context that is shared by all threads. Access to the context via its `ImplicitContext` interface is interlocked, so different threads can concurrently manipulate the context without risking data corruption or reading stale values.
- `PerThread`
  The communicator maintains a separate implicit context for each thread. This allows you to propagate contexts that depend on the sending thread (for example, to send per-thread transaction IDs).

See Also

- Explicit Request Contexts
- Per-Proxy Request Contexts

Copyright 2016, ZeroC, Inc.

Design Considerations for Request Contexts

On this page:

- Request Context Interactions
- Request Context Use Cases
- Recommendations for Request Contexts

Request Context Interactions

If you use explicit, per-proxy, and implicit contexts, it is important to be aware of their interactions:

- If you send an explicit context with an invocation, *only* that context is sent with the call, regardless of whether the proxy has a per-proxy context and whether the communicator has an implicit context.
- If you send an invocation via a proxy that has a per-proxy context, and the communicator also has an implicit context, the contents of the per-proxy and implicit context dictionaries are combined, so the *combination* of context entries of both contexts is transmitted to the server. If the per-proxy context and the implicit context contain the same key, but with different values, the *per-proxy* value takes precedence.

Request Context Use Cases

The purpose of `Ice::Context` is to permit services to be added to Ice that require some contextual information with every request. Contextual information can be used by services such as a transaction service (to provide the context of a currently established transaction) or a security service (to provide an authorization token to the server). IceStorm uses the context to provide an optional `cost` parameter to the service that influences how the service propagates messages to down-stream subscribers, and Glacier2 uses the context to influence request routing.

In general, services that require such contextual information can be implemented much more elegantly using contexts because this hides explicit Slice parameters that would otherwise have to be supplied by the application programmer with every call. For a request-forwarding intermediary service such as IceStorm or Glacier2, using a parameter is not an option because the service does not know the signatures of the requests that it is forwarding, but the service does have access to the context. For this use case, the context is a convenient way for the caller to annotate a request.

In addition, contexts, because they are optional, permit a single Slice definition to apply to implementations that use the context, as well as to implementations that do not use it. In this way, to add transactional semantics to an existing service, you do not need to modify the Slice definitions to add an extra parameter to each operation. (Adding an extra parameter would not only be inconvenient for clients, but would also split the type system into two halves: without contexts, we would need different Slice definitions for transactional and non-transactional implementations of (conceptually) a single service.)

Finally, per-proxy contexts permit context information to be passed through intermediate parts of your program without cooperation of those intermediate parts. For example, suppose you set a per-proxy context on a proxy and then pass that proxy to another function. When that function uses the proxy to invoke an operation, the per-proxy context will still be sent. In other words, per-proxy contexts allow you to transparently propagate information via intermediaries that are ignorant of the presence of any context.

Keep in mind though that this works only within a single process. If you stringify a proxy or transmit it as a parameter over the wire, the per-proxy context is *not* preserved. (Ice does not write the per-proxy context into stringified proxies and does not marshal the per-proxy context when a proxy is marshaled.)

Recommendations for Request Contexts

Contexts are a powerful mechanism for transparent propagation of context information, *if used correctly*. In particular, you may be tempted to use contexts as a means of versioning an application as it evolves over time. For example, version 2 of your application may accept two parameters on an operation that, in version 1, used to accept only a single parameter. Using contexts, you could supply the second parameter as a name-value pair to the server and avoid changing the Slice definition of the operation in order to maintain backward compatibility.

We *strongly* urge you to resist any temptation to use contexts in this manner. The strategy is fraught with problems:

- Missing context
  There is nothing that would compel a client to actually send a context when the server expects to receive a context: if a client forgets to send a context, the server, somehow, has to make do without it (or throw an exception).

- Missing or incorrect keys
  Even if the client does send a context, there is no guarantee that it has set the correct key. (For example, a simple spelling error can cause the client to send a value with the wrong key.)

- Incorrect values
  The value of a context is a string, but the application data that is to be sent might be a number, or it might be something more

Copyright 2016, ZeroC, Inc.

complex, such as a structure with several members. This forces you to encode the value into a string and decode the value again on the server side. Such parsing is tedious and error prone, and far less efficient than sending strongly-typed parameters. In addition, the server has to deal with string values that fail to decode correctly (for example, because of an encoding error made by the client).

None of the preceding problems can arise if you use proper Slice parameters: parameters cannot be accidentally omitted and they are strongly typed, making it much less likely for the client to accidentally send a meaningless value. Furthermore, as of Ice 3.5 you can use optional parameters to modify the signature of an operation without breaking backward compatibility. For a large-scale solution to application versioning, we recommend using Ice facets.

Contexts are meant to be used to transmit simple tokens (such as a transaction identifier) for services that cannot be reasonably implemented without them; you should restrict your use of contexts to that purpose and resist any temptation to use contexts for any other purpose.

Finally, be aware that, if a request is routed via one or more Ice routers, contexts may be dropped by intermediate routers if they consider them illegal. This means that, in general, you cannot rely on an arbitrary context value that is created by an application to actually still be present when a request arrives at the server — only those context values that are known to routers and that are considered legitimate are passed on. It follows that you should not abuse contexts to pass things that really should be passed as parameters.

See Also

- Explicit Request Contexts
- Per-Proxy Request Contexts
- Implicit Request Contexts
- Optional Values
- Versioning
- IceStorm

Copyright 2016, ZeroC, Inc.

Invocation Timeouts

On this page:

- Overview of Invocation Timeouts
- Configuring the Default Invocation Timeout
- Configuring Invocation Timeouts for Proxies
- Invocation Timeout Failures

Overview of Invocation Timeouts

Invocation timeouts let an application specify the maximum amount of time it's willing to wait for invocations to complete. If the timeout expires, the application receives `InvocationTimeoutException` as the result of an invocation.

Use connection timeouts to detect network failures in a reasonable period of time.

Invocation timeouts were introduced in Ice 3.6. In earlier Ice versions, connection timeouts were also used as invocation timeouts.

Configuring the Default Invocation Timeout

The property `Ice.Default.InvocationTimeout` establishes the default timeout value for invocations. This property has a default value of `-1`, which means invocations do not time out by default. If defined to `-2`, invocation timeouts are disabled and the Ice run time behaves like Ice versions < 3.6: it uses connection timeouts (defined on the endpoints) to wait for the response of the invocation. This setting is provided only for backward compatibility and might be deprecated in a future Ice major release.

Consider this setting:

```
Ice.Default.InvocationTimeout=5000
```

This configuration causes all invocations to time out if they do not complete within five seconds. Generally speaking however, it's unlikely that a single timeout value will be appropriate for all of the operations that an application invokes. It's more common for applications to configure invocation timeouts on a per-proxy basis, as we describe in the next section.

Configuring Invocation Timeouts for Proxies

You have a couple of options for configuring invocation timeouts at the proxy level:

- Use a proxy property
- Call `ice_invocationTimeout`

Assuming you've defined a configuration property containing a proxy that your application reads using `propertyToProxy`, you can statically configure an invocation timeout as follows:

```
# Assumes the application calls propertyToProxy("MyProxy")
MyProxy=theIdentity:tcp -p 5000
MyProxy.InvocationTimeout=2500  # 2.5 seconds
```

The `InvocationTimeout` proxy property specifies the invocation timeout that will be used for all invocations made via the proxy returned by `propertyToProxy`.

To configure an invocation timeout at run time, use the `ice_invocationTimeout` factory method to obtain a new proxy with the desired timeout:

Copyright 2016, ZeroC, Inc.

C++

```
Filesystem::FilePrx myFile = ...;
FileSystem::FilePrx timeoutFile = myFile->ice_invocationTimeout(2500);  // 2.5
seconds
```

Invocation Timeout Failures

An application that configures invocation timeouts must be prepared to catch `InvocationTimeoutException`:

C++

```
Filesystem::FilePrx myFile = ...;
FileSystem::FilePrx timeoutFile = myFile->ice_invocationTimeout(2500);

try {
    Lines text = timeoutFile->read();   // Read with timeout
} catch (const Ice::InvocationTimeoutException &) {
    cerr << "invocation timed out" << endl;
}

Lines text = myFile->read();            // Read without timeout
```

The effects of an invocation timeout are limited to the client; no indication is sent to the server, which may still be busy dispatching the request. The Ice run time in the client ignores a reply for this request if the server eventually sends one.

Ice does **not** perform automatic retries for invocation timeouts.

See Also

- Proxy Methods
- Proxy Properties
- Ice.Default.*
- Connection Timeouts
- Obtaining Proxies

Copyright 2016, ZeroC, Inc.

Automatic Retries

Ice may automatically retry a proxy invocation after a failure. This is a powerful feature that, when used in the proper situations, can significantly improve the robustness of your application without any additional programming effort. The retry facility is governed by one overriding principle: always respect at-most-once semantics. At-most-once semantics dictate that the Ice run time in the client must never retry a failed proxy invocation unless Ice guarantees that the server has not already received the request, or unless the application declares that it is safe for Ice to violate at-most-once semantics for the request.

To understand the importance of obeying at-most-once semantics, consider the following Slice definition:

Slice

```
interface Account {
    long withdraw(long amount);
};
```

The `withdraw` operation removes funds from an account. If an invocation of `withdraw` fails, automatically retrying the request introduces the risk of a duplicate withdrawal unless Ice is absolutely sure that the server has not already executed the request.

This page examines automatic retries in more detail.

On this page:

- Automatic Retries for Request Failures
- Automatic Retries for Idempotent Operations
- Configuring Automatic Retries
  - Retry Intervals
  - Retry Logging
- Timeouts and Automatic Retries
- Connections and Automatic Retries
  - Connection Errors
  - Connection Status

Automatic Retries for Request Failures

Ice considers a request to have failed if any of the following conditions are true:

- A connection could not be established
- A connection was lost before the reply was received
- A timeout expired
- An exception occurred while sending the request or receiving the reply
- An error occurred in the server while dispatching the request that causes the server to return an `UnknownException` or `RequestFailedException`

Ice considers an invocation that results in a user exception to be successful and therefore excludes it from consideration for automatic retries.

Ice must determine the answers to several questions to decide whether to retry a failed request:

1. What kind of error caused the request to fail?

Ice does not bother retrying a request if it knows the same error is going to occur again. For example, Ice never retries an invocation that raises a `MarshalException`, which indicates that there was a problem while encoding or decoding a message. Retrying such an invocation is unlikely to change the outcome.

Ice also never retries exceptions that derive from `RequestFailedException` because they indicate a permanent failure. One such subclass is `OperationNotExistException`, whose occurrence signals a serious problem in the application. For instance, it might mean that the client and server are using incompatible Slice definitions, or that the client is trying to invoke operations on the wrong object. The exception to this rule is `ObjectNotExistException`, which Ice does consider to be worthy of retry if the proxy in question is indirect beca use it gives an application the ability to transparently migrate an Ice object.

In addition to user exceptions and subclasses of `RequestFailedException`, a server can also return an instance of `UnknownException` , `UnknownLocalException`, or `UnknownUserException` to indicate that it encountered an unexpected exception while dispatching the

Copyright 2016, ZeroC, Inc.

request. These exceptions *are* eligible for retry.

2. When did the error occur?

If the error is still a candidate for retry, Ice needs to know whether the server has received the request. Naturally, the Ice run time in the client cannot possibly know that information until the server confirms it by sending a reply. However, to be conservative Ice assumes that the server has received the request as soon as Ice has written the *entire* protocol message to the client's local transport buffers. If the error occurred before Ice managed to write the complete message, retrying the request would not violate at-most-once semantics.

The Ice run time in the server also has the ability to notify the client that a request was not dispatched and therefore that it is safe for the Ice run time in the client to retry the request without violating at-most-once semantics. For example, this situation can occur when the server is shutting down while there are pending requests that have yet to be executed. Sending this notification allows a client to transparently fail over to another server.

3. Does the application require strict adherence to at-most-once semantics for this request?

An application can grant permission for Ice to violate at-most-once semantics for certain Slice operations by marking them as *idempotent*, causing Ice to retry a request that otherwise would be ineligible because the server has already received it. We discuss idempotent operations in more detail below.

If Ice determines that an invocation cannot be retried, it raises the exception that caused the request failure to the application. On the other hand, if Ice does retry the invocation and the subsequent retries also fail, Ice raises the *last* exception to the application. For example, if the first attempt fails with `ConnectionRefusedException` and the retry fails with `ConnectTimeoutException`, the invocation raises `ConnectTimeoutException` to the application.

Automatic Retries for Idempotent Operations

Annotating a Slice operation with the `idempotent` keyword notifies Ice that it can safely violate at-most-once semantics:

Slice

```
interface Account {
    long withdraw(long amount);
    idempotent long getBalance();
};
```

Although `withdraw` clearly requires the stricter treatment, there is no harm in automatically retrying the `getBalance` operation even if the server executes the same request more than once.

In general, "read-only" operations are good candidates for the `idempotent` keyword whereas many mutating operations are not. However, the risk of duplicate requests is acceptable even for some kinds of mutating operations:

Slice

```
interface Account {
    long withdraw(long amount);
    idempotent long getBalance();
    idempotent void changeAddress(string newAddress);
};
```

Here we have marked `changeAddress` as idempotent because executing the request twice has the same effect as executing it only once.

The benefit of the `idempotent` keyword and the associated relaxation of retry semantics is that an invocation that otherwise might have raised an exception has at least one more chance to succeed. Furthermore, the application does not need to initiate the retry, and in fact the retry activities are completely transparent: if a subsequent retry succeeds, the application receives its results as if nothing went wrong. The invocation only raises an exception once Ice has reached its configured retry limits.

Configuring Automatic Retries

Copyright 2016, ZeroC, Inc.

Retry Intervals

The `Ice.RetryIntervals` property configures the retry behavior for a communicator and affects invocations on every proxy created by that communicator. (Retry behavior cannot be configured on a per-proxy basis.) The value of this property consists of a series of integers separated by whitespace. The number of integers determines how many retry attempts Ice makes, and the value of each entry represents a delay in milliseconds. If this property is not defined, the default behavior is to retry once immediately after the first failure, which is equivalent to the following property definition:

```
Ice.RetryIntervals=0
```

You may want a more elaborate configuration for your application, such as a gradual increase in the delay between retries:

```
Ice.RetryIntervals=0 100 500 1000
```

With this setting, Ice retries immediately as in the default case. If the first retry attempt also fails, Ice waits 100 milliseconds before trying again, then 500 milliseconds, and finally tries one more time after waiting one second.

In some situations you may need to disable retries completely. For example, an application might implement its own retry logic and therefore require immediate notification when a failure occurs. Clients that establish a session with a Glacier2 router also need to disable retries. To prevent automatic retries, use a value of `-1`:

```
Ice.RetryIntervals=-1
```

Retry Logging

To monitor Ice's retry activities, configure your program with the property `Ice.Trace.Retry` set to a non-zero value:

```
Ice.Trace.Retry=1
```

When retry tracing is enabled, Ice logs a message each time it attempts a retry; the log message includes a description of the exception that prompted the retry. Ice also logs a message when it reaches the retry limit.

You can configure Ice to log even more information about retries by setting the property to `2`:

```
Ice.Trace.Retry=2
```

This setting prompts Ice to include additional details about connections and endpoints.

Timeouts and Automatic Retries

If a proxy invocation fails due to an invocation timeout or connection timeout, the application must be prepared for Ice to raise an `InvocationTimeoutException` or `TimeoutException`, respectively. However, a developer that is testing timeouts in an application may be initially confused to discover that it's taking twice as long as expected for Ice to raise the exception. Automatic retries are usually the reason for this situation.

For example, suppose that a proxy is configured with a one-second invocation timeout and automatic retries are enabled with the default setting. If an invocation on that proxy fails due to a timeout and Ice determines that the invocation is eligible for retry (using the criteria described above), Ice immediately tries the invocation again and waits for another timeout period to expire before finally raising `InvocationTimeoutException`. From the application's perspective, the invocation fails after approximately two seconds.

Consequently, you can compute an approximate worst-case timeout value as follows:

Copyright 2016, ZeroC, Inc.

```
T = t * (N + 1) + D
```

where *t* is the timeout value, *N* is the number of retry intervals, and *D* is the sum of the retry intervals (the total delay between retries). Consider our example again:

```
Ice.RetryIntervals=0 100 500 1000
```

Using this configuration with a one-second timeout, our approximate worst-case timeout is `1 * 5 + 1.6 = 6.6` seconds.

Connections and Automatic Retries

The behavior of automatic retries is intimately tied to the presence (and absence) of connections. This section describes the errors that cause Ice to close connections, and provides more details about how connections influence retries.

Connection Errors

Ice automatically closes a connection in response to certain fatal error conditions. Of these, the one that is the most likely to affect Ice applications is a connection timeout. Other errors that prompt Ice to close a connection include the following:

- a socket failure while performing I/O on the connection
- receiving an improperly formatted message
- dispatching an operation to a Java servant raises `OutOfMemoryError` or `AssertionError`

When Ice closes a connection in response to one of these errors, all other outstanding requests on the same connection also fail and may be retried if eligible.

Connection Status

One factor that influences retry behavior is the status of the connection on which the failed request was attempted. If the failure caused Ice to close the connection (as discussed in the previous section), or if the request failed because Ice could not establish a connection, Ice must try to obtain another connection before it can retry the request.

It is also important to understand that Ice may not retry the invocation on the original endpoint *even if the connection that was used for the initial request remains open*. The retry behavior in this case depends on several criteria:

- whether the proxy caches its connection
- whether the proxy contains multiple endpoints
- whether other connections exist to any of the proxy's endpoints
- the proxy's configured endpoint selection type

Generally speaking, you must configure your application carefully if you need fine-grained control over Ice's retry behavior.

See Also

- Terminology
- Operations
- Invocation Timeouts
- Connection Timeouts
- Connection Establishment
- Getting Started with Glacier2

Copyright 2016, ZeroC, Inc.

Oneway Invocations

On this page:

- Design Considerations for Oneway Invocations
- Creating Oneway Proxies

Design Considerations for Oneway Invocations

A oneway invocation is sent on the client side by writing the request to the client's local transport buffers; the invocation completes and returns control to the application code as soon as it has been accepted by the local transport. Of course, this means that a oneway invocation is unreliable: it may never be sent (for example, because of a network failure) or it may not be accepted in the server (for example, because the target object does not exist).

This is an issue in particular if you use Active Connection Management (ACM): if a server closes a connection at the wrong moment, it's possible for the client to lose already-buffered oneway requests. If your client uses oneway (or batched oneway) requests, we recommend that you either disable ACM for the server side, or enable ACM heartbeats in the client to ensure the connection remains active. In addition, if clients use oneway requests and your application initiates server shutdown, it is the responsibility of your application to ensure either that it can cope with the potential loss of buffered oneway requests, or that it does not shut down the server at the wrong moment (while clients still have oneway requests that are buffered, but not yet sent).

If anything goes wrong with a oneway request, the client-side application code does not receive any notification of the failure; the only errors that are reported to the client are local errors that occur on the client side during call invocation (such as failure to establish a connection, for example).

As a consequence of oneway invocation, if you call `ice_ping` on a oneway proxy, successful completion does *not* indicate that the target object exists and could successfully be contacted — you will receive an exception only if something goes wrong on the client side, but not if something goes wrong on the server side. Therefore, if you want to use `ice_ping` with a oneway proxy and be certain that the target object exists and can successfully be contacted, you must first convert the oneway proxy into a twoway proxy. For example, in C++:

C++

```
SomeObjectPrx onewayPrx = ...; // Get a oneway proxy

try {
    onewayPrx->ice_twoway()->ice_ping();
} catch(const Ice::Exception&)
    cerr << "object not reachable" << endl;
}
```

Oneway invocations are received and processed on the server side like any other incoming request. If necessary, a server can distinguish a oneway invocation by examining the `requestId` member of `Ice::Current`: a non-zero value denotes a twoway request, whereas a value of zero indicates a oneway request.

Oneway invocations do not incur any return traffic from the server to the client: the server never sends a reply message in response to a oneway invocation. This means that oneway invocations can result in large efficiency gains, especially for large numbers of small messages, because the client does not have to wait for the reply to each message to arrive before it can send the next message.

In order to be able to invoke an operation as oneway, two conditions must be met:

- The operation must have a `void` return type, must not have any out-parameters, and must not have an exception specification.

  This requirement reflects the fact that the server does not send a reply for a oneway invocation to the client: without such a reply, there is no way to return any values or exceptions to the client. If you attempt to invoke an operation that returns values to the client as a oneway operation, the Ice run time throws a `TwowayOnlyException`.

- The proxy on which the operation is invoked must support a stream-oriented transport (such as TCP or SSL).

  Oneway invocations require a stream-oriented transport. (To get something like a oneway invocation for datagram transports, you need to use a datagram invocation.) If you attempt to create a oneway proxy for an object that does not offer a stream-oriented transport, the Ice run time throws a `NoEndpointException`.

Despite their theoretical unreliablity, in practice, oneway invocations are reliable (but not infallible [1]): they are sent via a stream-oriented

Copyright 2016, ZeroC, Inc.

transport, so they cannot get lost except when the connection is shutting down or fails entirely. In particular, the transport uses its usual flow control, so the client cannot overrun the server with messages. On the client-side, the Ice run time will block if the client's transport buffers fill up, so the client-side application code cannot overrun its local transport.

Consequently, oneway invocations normally do not block the client-side application code and return immediately, provided that the client does not consistently generate messages faster than the server can process them. If the rate at which the client invokes operations exceeds the rate at which the server can process them, the client-side application code will eventually block in an operation invocation until sufficient room is available in the client's transport buffers to accept the invocation. If your application requires that oneway requests never block the calling thread, you can use asynchronous oneway invocations instead.

Regardless of whether the client exceeds the rate at which the server can process incoming oneway invocations, the execution of oneway invocations in the server proceeds asynchronously: the client's invocation completes before the message even arrives at the server.

One thing you need to keep in mind about oneway invocations is that they may appear to be reordered in the server: because oneway invocations are sent via a stream-oriented transport, they are guaranteed to be received in the order in which they were sent. However, the server's thread pool may dispatch each invocation in its own thread; because threads are scheduled preemptively, this may cause an invocation sent later by the client to be dispatched and executed before an invocation that was sent earlier. If oneway requests must be dispatched in order, you can use one of the serialization techniques described in Thread Pool Design Considerations.

For these reasons, oneway invocations are usually best suited to simple updates that are otherwise stateless (that is, do not depend on the surrounding context or the state established by previous invocations).

Creating Oneway Proxies

Ice selects between twoway, oneway, and datagram invocations via the proxy that is used to invoke the operation. By default, all proxies are created as twoway proxies. To invoke an operation as oneway, you must create a new proxy configured specifically for oneway invocations. The `ice_oneway` factory method is provided for this purpose. The Slice definition of `ice_oneway` would look as follows:

Slice

```
Object* ice_oneway();
```

We can call `ice_oneway` to create a oneway proxy and then use the proxy to invoke an operation as follows:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx o = communicator->stringToProxy(/* ... */);

// Get a oneway proxy.
//
Ice::ObjectPrx oneway;
try {
    oneway = o->ice_oneway();
} catch (const Ice::NoEndpointException&) {
    cerr << "No endpoint for oneway invocations" << endl;
}

// Down-cast to actual type.
//
PersonPrx onewayPerson = PersonPrx::uncheckedCast(oneway);

// Invoke an operation as oneway.
//
try {
    onewayPerson->someOp();
} catch (const Ice::TwowayOnlyException&) {
    cerr << "someOp() is not oneway" << endl;
}
```

Note that we use an `uncheckedCast` to down-cast the proxy from `ObjectPrx` to `PersonPrx`: for a oneway proxy, we cannot use a `checkedCast` because a `checkedCast` requires a reply from the server but, of course, a oneway proxy does not permit that reply. If instead you want to use a safe down-cast, you can first down-cast the twoway proxy to the actual object type and then obtain the oneway proxy:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx o = communicator->stringToProxy(/* ... */);

// Safe down-cast to actual type.
//
PersonPrx person = PersonPrx::checkedCast(o);

if (person) {
    // Get a oneway proxy.
    //
    PersonPrx onewayPerson;
    try {
        onewayPerson = person->ice_oneway();
    } catch (const Ice::NoEndpointException&) {
        cerr << "No endpoint for oneway invocations" << endl;
    }

    // Invoke an operation as oneway.
    //
    try {
        onewayPerson->someOp();
    } catch (const Ice::TwowayOnlyException&) {
        cerr << "someOp() is not oneway" << endl;
    }
}
```

Note that, while the second version of this code is somewhat safer (because it uses a safe down-cast), it is also slower (because the safe down-cast incurs the cost of an additional twoway message).

See Also

- Terminology
- Active Connection Management
- Batched Invocations
- The Current Object
- The Ice Threading Model
- Datagram Invocations
- The Ice Protocol

References

1. Snader, J. C. 2000. Effective TCP/IP Programming. Reading, MA: Addison-Wesley.

Copyright 2016, ZeroC, Inc.

Datagram Invocations

On this page:

- Design Considerations for Datagram Invocations
- Creating Datagram Proxies
- Using UDP Multicast

Design Considerations for Datagram Invocations

Datagram invocations are the equivalent of oneway invocations for datagram transports. As for oneway invocations, datagram invocations can be sent only for operations that have a `void` return type and do not have out-parameters or an exception specification. Attempts to use a datagram invocation with an operation that does not meet these criteria result in a `TwowayOnlyException`. In addition, datagram invocations can only be used if the proxy's endpoints include at least one UDP transport; otherwise, the Ice run time throws a `NoEndpointException`.

The semantics of datagram invocations are similar to oneway invocations: no return traffic flows from the server to the client and proceed asynchronously with respect to the client; a datagram invocation completes as soon as the client's transport has accepted the invocation into its buffers. However, datagram invocations differ in one respect from oneway invocations in that datagram invocations optionally support multicast semantics. Furthermore, datagram invocations have additional error semantics:

- Individual invocations may be lost or received out of order.

  On the wire, datagram invocations are sent as true datagrams, that is, individual datagrams may be lost, or arrive at the server out of order. As a result, not only may operations be dispatched out of order, an individual invocation out of a series of invocations may be lost. (This cannot happen for oneway invocations because, if a connection fails, *all* invocations are lost once the connection breaks down.)

- UDP packets may be duplicated by the transport.

  Because of the nature of UDP routing, it is possible for datagrams to arrive in duplicate at the server. This means that, for datagram invocations, Ice does *not* guarantee at-most-once semantics: if UDP datagrams are duplicated, the same invocation may be dispatched more than once in the server.

- UDP packets are limited in size.

  The maximum size of an IP datagram is 65,535 bytes. Of that, the IP header consumes 20 bytes, and the UDP header consumes 8 bytes, leaving 65,507 bytes as the maximum payload. If the marshaled form of an invocation, including the Ice request header exceeds that size, the invocation is lost. (Exceeding the size limit for a UDP datagram is indicated to the application by a `DatagramLimitException`.)

Because of their unreliable nature, datagram invocations are best suited to simple update messages that are otherwise stateless. In addition, due to the high probability of loss of datagram invocations over wide area networks, you should restrict use of datagram invocations to local area networks, where they are less likely to be lost. (Of course, regardless of the probability of loss, you must design your application such that it can tolerate lost or duplicated messages.)

Creating Datagram Proxies

To invoke an operation as datagram, you must create a new proxy configured specifically for datagram invocations. The `ice_datagram` factory method is provided for this purpose. The Slice definition of `ice_datagram` would look as follows:

Slice

```
Object* ice_datagram();
```

We can call `ice_datagram` to create a oneway proxy and then use the proxy to invoke an operation as follows:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx o = communicator->stringToProxy(/* ... */);

// Get a datagram proxy.
//
Ice::ObjectPrx datagram;
try {
    datagram = o->ice_datagram();
} catch (const Ice::NoEndpointException&) {
    cerr << "No endpoint for datagram invocations" << endl;
}

// Down-cast to actual type.
//
PersonPrx datagramPerson = PersonPrx::uncheckedCast(datagram);

// Invoke an operation as a datagram.
//
try {
    datagramPerson->someOp();
} catch (const Ice::TwowayOnlyException&) {
    cerr << "someOp() is not oneway" << endl;
}
```

As for the oneway example, you can alternatively choose to first do a safe down-cast to the actual type of interface and then obtain the datagram proxy, rather than relying on an unsafe down-cast, as shown above. However, doing so may be disadvantageous for two reasons:

- Safe down-casts are sent via a stream-oriented transport. This means that using a safe down-cast will result in opening a connection for the sole purpose of verifying that the target object has the correct type. This is expensive if all the other traffic to the object is sent via datagrams.
- If the proxy does not offer a stream-oriented transport, the `checkedCast` fails with a `NoEndpointException`, so you can use this approach only for proxies that offer both a UDP endpoint and a TCP/IP and/or SSL endpoint.

Using UDP Multicast

The UDP transport included in Ice for C++, Java and .NET also supports IP multicast. Assuming it's enabled on your host, using IP multicast in your application can be as simple as changing the host in the UDP endpoint to an IPv4 or IPv6 address in the multicast range:

```
# Object Adapter endpoint:
Discover.Endpoints=udp -h 239.255.1.1 -p 10000

# Corresponding proxy endpoint:
Discover.Proxy=discover:udp -h 239.255.1.1 -p 10000
```

You can optionally select the network interface to use for multicast endpoints by including the `--interface` option.

In one respect, using multicast in Ice is no different than using regular datagram invocations; all of the design considerations mentioned above still apply. However, the fact that there could be any number of listeners (or none at all) adds new possibilities for your application design. The Ice distribution includes a simple example of a multicast application in `demo/Ice/multicast`.

Consider using the IceDiscovery plug-in if your objective in using multicast is the discovery of available servers.

Copyright 2016, ZeroC, Inc.

See Also

- Terminology
- Oneway Invocations
- The Ice Protocol

Copyright 2016, ZeroC, Inc.

Batched Invocations

Oneway and datagram invocations are normally sent as individual messages, that is, the Ice run time sends the oneway or datagram invocation to the server immediately, as soon as the client makes the call. If a client sends a number of oneway or datagram invocations in succession, the client-side run time traps into the OS kernel for each message, which is expensive. In addition, each message is sent with its own message header, that is, for *N* messages, the bandwidth for *N* message headers is consumed. In situations where a client sends a number of oneway or datagram invocations, the additional overhead can be considerable.

To avoid the overhead of sending many small messages, you can send oneway and datagram invocations in a batch: instead of being sent as a separate message, a batch invocation is placed into a client-side buffer by the Ice run time. Successive batch invocations are added to the buffer and accumulated on the client side until they are flushed, either explicitly by the client or automatically by the Ice run time.

On this page:

- Proxy Methods for Batched Invocations
- Automatically Flushing Batched Invocations
- Batched Invocations for Fixed Proxies
- Considerations for Batched Datagrams
- Compressing Batched Invocations
- Active Connection Management and Batched Invocations
- Batched Invocation Interceptors

Proxy Methods for Batched Invocations

Several proxy methods support the use of batched invocations. In Slice, these methods would look as follows:

Slice

```
Object* ice_batchOneway();
Object* ice_batchDatagram();
void ice_flushBatchRequests();
```

The `ice_batchOneway` and `ice_batchDatagram` methods create a new proxy configured for batch invocations. Once you obtain a batch proxy, messages sent via that proxy are buffered by the proxy instead of being sent immediately. Once the client has invoked one or more operations on a batch proxy, it can call `ice_flushBatchRequests` to explicitly flush the batched invocations. This causes the batched messages to be sent "in bulk", preceded by a single message header. On the server side, batched messages are dispatched by a single thread, in the order in which they were written into the batch. This means that messages from a single batch cannot appear to be reordered in the server. Moreover, either all messages in a batch are delivered or none of them. (This is true even for batched datagrams.)

Asynchronous versions of `ice_flushBatchRequests` are also available; see the relevant language mapping for more information.

As of Ice 3.6, batched invocations are queued by the proxy on which the request was invoked (this is true for all proxies except fixed proxies) . It's important to be aware of this behavior for several reasons:

- Batched invocations queued on a proxy will be lost if that proxy is deallocated prior to being flushed
- Proxy instances maintain separate queues even if they refer to the same target object
- Using proxy factory methods may (or may not) create new proxy instances, which affects how batched invocations are queued. Consider this example:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::CommunicatorPtr comm = ...;
Ice::ObjectPrx batch = comm->stringToProxy("...")->ice_batchOneway(); //
Creates a batch oneway proxy
Ice::ObjectPrx secureBatch = batch->ice_secure(); // Might create a new proxy
instance
batch->ice_ping();
secureBatch->ice_ping();
batch->ice_flushBatchRequests(); // Might also flush requests on secureBatch
```

At run time, the `batch` and `secureBatch` variables might refer to the same proxy instance or two separate proxy instances, depending on the original stringified proxy's configuration. As a result, flushing requests using one variable may or may not flush the requests of the other.

Calling `ice_flushBatchRequests` on a proxy behaves like a oneway invocation in that automatic retries take place and an exception is raised if an error occurs while establishing a connection or sending the batch message.

Automatically Flushing Batched Invocations

The default behavior of the Ice run time, as governed by the configuration property `Ice.BatchAutoFlushSize`, automatically flushes batched invocations as soon as a batched request causes the accumulated message to exceed the specified limit. When this occurs, the Ice run time immediately flushes the existing batch of requests and begins a new batch with this latest request as its first element.

For batched oneway invocations, the value of `Ice.BatchAutoFlushSize` specifies the maximum message size in kilobytes; the default value is 1MB. In the case of batched datagram invocations, the maximum message size is the smaller of the system's maximum size for datagram packets and the value of `Ice.BatchAutoFlushSize`.

The receiver's setting for `Ice.MessageSizeMax` determines the maximum size that the Ice run time will accept for an incoming protocol message. The sender's setting for `Ice.BatchAutoFlushSize` must not exceed this limit, otherwise the receiver will silently discard the entire batch.

Automatic flushing is enabled by default as a convenience for clients to ensure a batch never exceeds the configured limit. A client can track batch request activity, and even implement its own auto-flush logic, by installing an interceptor.

Batched Invocations for Fixed Proxies

A *fixed proxy* is a special form of proxy that an application explicitly creates for use with a specific bidirectional connection. Batched invocations on a fixed proxy are not queued by the proxy, as is the case for regular proxies, but rather by the connection associated with the fixed proxy. Automatic flushing continues to work as usual for batched invocations on fixed proxies, and you have three options for manually flushing:

- Calling `ice_flushBatchRequests` on a fixed proxy flushes all batched requests queued by its connection; this includes batched requests from other fixed proxies that share the same connection
- Calling `Connection::flushBatchRequests` flushes all batched requests queued by the target connection
- Calling `Communicator::flushBatchRequests` flushes all batched requests on all connections associated with the target communicator

As of Ice 3.6, `Connection::flushBatchRequests` and `Communicator::flushBatchRequests` have no effect on batched requests queued by regular (non-fixed) proxies.

The synchronous versions of `flushBatchRequests` block the calling thread until the batched requests have been successfully written to the local transport. To avoid the risk of blocking, you must use the asynchronous versions instead (assuming they are supported by your chosen language mapping).

Note the following limitations in case a connection error occurs:

Copyright 2016, ZeroC, Inc.

- Any requests queued by that connection are lost
- Automatic retries are not attempted
- The proxy method `ice_flushBatchRequests` and `Connection::flushBatchRequests` will raise an exception, but `Communicator::flushBatchRequests` ignores any errors

`Connection::flushBatchRequests` and `Communicator::flushBatchRequests` always send the batched requests without compression. To send a batch of fixed-proxies request with compression, you need to call `ice_flushBatchRequests` on a fixed proxy created with `ice_compress(true)`, for example:

C++

```
fixedPrx->ice_compress(true)->ice_flushBatchRequests();
```

Considerations for Batched Datagrams

For batched datagram invocations, you need to keep in mind that, if the data for the invocations in a batch substantially exceeds the PDU size of the network, it becomes increasingly likely for an individual UDP packet to get lost due to fragmentation. In turn, loss of even a single packet causes the entire batch to be lost. For this reason, batched datagram invocations are most suitable for simple interfaces with a number of operations that each set an attribute of the target object (or interfaces with similar semantics). Batched oneway invocations do not suffer from this risk because they are sent over stream-oriented transports, so individual packets cannot be lost.

If automatic flushing is enabled, Ice's default behavior uses the smaller of `Ice.BatchAutoFlushSize` and `Ice.UDP.SndSize` to determine the maximum size for a batch datagram message.

Compressing Batched Invocations

Batched invocations are more efficient if you also enable compression for the transport: many isolated and small messages are unlikely to compress well, whereas batched messages are likely to provide better compression because the compression algorithm has more data to work with.

Regardless of whether you used batched messages or not, you should enable compression only on lower-speed links. For high-speed LAN connections, the CPU time spent doing the compression and decompression is typically longer than the time it takes to just transmit the uncompressed data.

Active Connection Management and Batched Invocations

As for oneway invocations, server-side Active Connection Management (ACM) can interfere with batched invocations over TCP or TCP-based transports (SSL, WebSocket, etc.). With server-side ACM enabled, it's possible for a server to close the connection at the wrong moment and not process a batch – with no indication being returned to the client that the batch was lost. We recommend that you either disable ACM for the server side, or enable ACM heartbeats in the client to ensure the connection remains active.

Batched Invocation Interceptors

Ice 3.6 added a batch interceptor API that allows you to implement your own auto-flush algorithm or receive notification when an auto-flush fails. We present the C++ API below (the API for other languages is very similar):

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {
class BatchRequest {
public:
    virtual void enqueue() const = 0;
    virtual int getSize() const = 0;
    virtual const std::string& getOperation() const = 0;
    virtual const Ice::ObjectPrx& getProxy() const = 0;
};

class BatchRequestInterceptor : public IceUtil::Shared {
public:
    virtual void enqueue(const BatchRequest& req, int count, int size) = 0;
};
typedef IceUtil::Handle<BatchRequestInterceptor> BatchRequestInterceptorPtr;

// C++11 only:
BatchRequestInterceptorPtr
newBatchRequestInterceptor(const std::function<void (const BatchRequest&, int,
int)>&);
}
```

You install an interceptor by setting the `batchRequestInterceptor` member of the `InitializationData` object that the application constructs when initializing a new communicator:

C++

```
namespace Ice {
struct InitializationData {
    ...
    BatchRequestInterceptorPtr batchRequestInterceptor;
};
}
```

The Ice run time invokes the `enqueue` method on the interceptor for each batch request, passing the following arguments:

- `req` - An object representing the batch request being queued
- `count` - The number of requests currently in the queue
- `size` - The number of bytes consumed by the requests currently in the queue

The request represented by `req` is not included in the `count` and `size` figures.

A batch request is not queued until the interceptor calls `BatchRequest::enqueue`. The minimal interceptor implementation is therefore:

Copyright 2016, ZeroC, Inc.

C++

```
class BatchRequestInterceptorI : public Ice::BatchRequestInterceptor {
public:
    virtual void enqueue(const BatchRequest& req, int count, int size)
    {
        req.enqueue();
    }
};
```

A more sophisticated implementation might use its own logic for automatically flushing queued requests:

C++

```
class BatchRequestInterceptorI : public Ice::BatchRequestInterceptor {
public:
    BatchRequestInterceptorI(const Ice::CommunicatorPtr& communicator)
    {
        _limit =
communicator->getProperties()->getPropertyAsIntWithDefault("Ice.BatchAutoFlushSize"
, 1024);
        _limit *= 1024; // Convert to bytes
    }

    virtual void enqueue(const BatchRequest& req, int count, int size)
    {
        if(size + req.getSize() > _limit)
            req.getProxy()->begin_ice_flushBatchRequests(
                [=](const Ice::Exception& ex) { /* Handle error */ },
                [=](bool sentSynchronously) { /* Sent callback */ });

        req.enqueue();
    }

private:
    int _limit;
};
```

In this example, the implementation consults the existing Ice property `Ice.BatchAutoFlushSize` to determine the limit that triggers an automatic flush. If a flush is necessary, the interceptor can obtain the relevant proxy by calling `getProxy` on the `BatchRequest` object.

Specifying your own exception handler when calling `begin_ice_flushBatchRequests` gives you the ability to take action if a failure occurs (Ice's default automatic flushing implementation ignores any errors). Aside from logging a message, your options are somewhat limited because it's not possible for the interceptor to force a retry.

For datagram proxies, we strongly recommend using a maximum queue size that is smaller than the network MTU to minimize the risk that datagram fragmentation could cause an entire batch to be lost.

See Also

- Oneway Invocations

Copyright 2016, ZeroC, Inc.

- Datagram Invocations
- Communicators
- Using Connections
- The Ice Protocol
- Protocol Compression
- Active Connection Management

Copyright 2016, ZeroC, Inc.

Server-Side Features

This section presents all APIs and features that are specific to server applications.

Client-specific features are presented in Client-Side Features, while transverse features best understood while considering both sides of a client-server interaction are described in Client-Server Features.

Topics

- Object Adapters
- The Current Object
- Servant Locators
- Default Servants
- Dispatch Interceptors

Copyright 2016, ZeroC, Inc.

Object Adapters

A communicator contains one or more object adapters. An object adapter sits at the boundary between the Ice run time and the server application code and has a number of responsibilities:

- Maps Ice objects to servants for incoming requests and dispatches the requests to the application code in each servant (that is, an object adapter implements an up-call interface that connects the Ice run time and the application code in the server).
- Assists in life cycle operations so Ice objects and servants can be created and existing destroyed without race conditions.
- Provides one or more transport endpoints. Clients access the Ice objects provided by the adapter via those endpoints. (It is also possible to create an object adapter without endpoints. In this case the adapter is used for bidirectional callbacks.

Each object adapter has one or more servants that incarnate Ice objects, as well as one or more transport endpoints. If an object adapter has more than one endpoint, all servants registered with that adapter respond to incoming requests on any of the endpoints. In other words, if an object adapter has multiple transport endpoints, those endpoints represent alternative communication paths to the same set of objects (for example, via different transports).

Each object adapter belongs to exactly one communicator (but a single communicator can have many object adapters). Each object adapter has a name that distinguishes it from all other object adapters in the same communicator.

Each object adapter can optionally have its own thread pool, enabled via the `<adapter-name>.ThreadPool.Size` property. If so, client invocations for that adapter are dispatched in a thread taken from the adapter's thread pool instead of using a thread from the communicator's server thread pool.

Servants are the physical manifestation of an Ice object, that is, they are entities that are implemented in a concrete programming language and instantiated in the server's address space. Servants provide the server-side behavior for operation invocations sent by clients. The same servant can be registered with one or more object adapters.

Topics

- The Active Servant Map
- Creating an Object Adapter
- Servant Activation and Deactivation
- Object Adapter States
- Object Adapter Endpoints
- Creating Proxies with an Object Adapter
- Using Multiple Object Adapters

See Also

- Object Adapter Thread Pools
- Bidirectional Connections

Copyright 2016, ZeroC, Inc.

The Active Servant Map

Each object adapter maintains a data structure known as the active servant map.

On this page:

- Role of the Active Servant Map
- Design Considerations for the Active Servant Map

Role of the Active Servant Map

The *active servant map* (or *ASM*, for short) is a lookup table that maps object identities to servants: for C++, the lookup value is a smart pointer to the corresponding servant's location in memory; for Java and C#, the lookup value is a reference to the servant. When a client sends an operation invocation to the server, the request is targeted at a specific transport endpoint. Implicitly, the transport endpoint identifies the object adapter that is the target of the request (because no two object adapters can be bound to the same endpoint). The proxy via which the client sends its request contains the object identity for the corresponding object, and the client-side run time sends this object identity over the wire with the invocation. In turn, the object adapter uses that object identity to look in its ASM for the correct servant to dispatch the call to, as shown below:

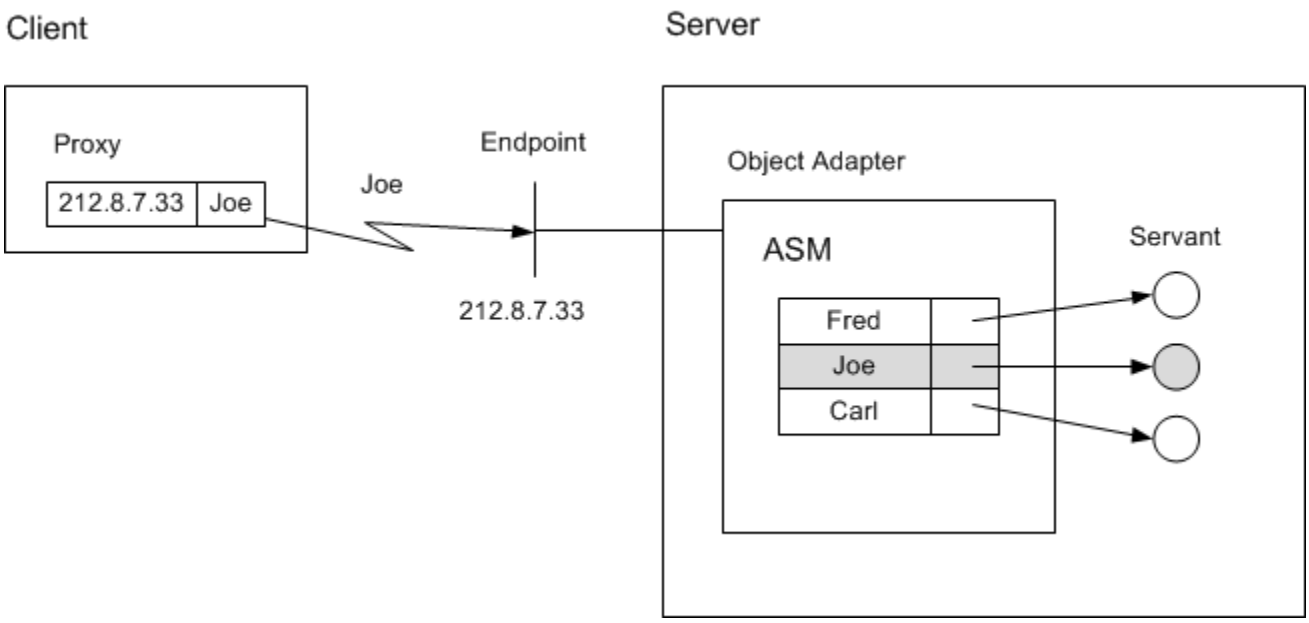

Binding a request to the correct servant.

The process of associating a request via a proxy to the correct servant is known as *binding*. The scenario depicted in the illustration shows direct binding, in which the transport endpoint is embedded in the proxy. Ice also supports an indirect binding mode, in which the correct transport endpoints are provided by a location service.

If a client request contains an object identity for which there is no entry in the adapter's ASM, the adapter returns an `ObjectNotExistException` to the client (unless you use a default servant or servant locator).

Design Considerations for the Active Servant Map

Using an adapter's ASM to map Ice objects to servants has a number of design implications:

- Each Ice object is represented by a different servant.

  > It is possible to register a single servant with multiple identities. However, there is little point in doing so because a default servant achieves the same thing.

- All servants for all Ice objects are permanently in memory.

Using a separate servant for each Ice object in this fashion is common to many server implementations: the technique is simple to implement and provides a natural mapping from Ice objects to servants. Typically, on start-up, the server instantiates a separate servant for each Ice object, activates each servant, and then calls `activate` on the object adapter to start the flow of requests.

Copyright 2016, ZeroC, Inc.

There is nothing wrong with the above design, provided that two criteria are met:

1. The server has sufficient memory available to keep a separate servant instantiated for each Ice object at all times.
2. The time required to initialize all the servants on start-up is acceptable.

For many servers, neither criterion presents a problem: provided that the number of servants is small enough and that the servants can be initialized quickly, this is a perfectly acceptable design. However, the design does not scale well: the memory requirements of the server grow linearly with the number of Ice objects so, if the number of objects gets too large (or if each servant stores too much state), the server runs out of memory.

Ice offers two APIs that help you scale servers to larger numbers of objects: *servant locators* and *default servants*. A default servant is essentially a simplified version of a servant locator that satisfies the majority of use cases, whereas a servant locator provides more flexibility for those applications that require it.

See Also

- Servant Locators
- Default Servants
- Locators

Copyright 2016, ZeroC, Inc.

Creating an Object Adapter

An object adapter is an instance of the local interface `ObjectAdapter`. You create an object adapter by calling one of several operations on a communicator:

Slice

```
module Ice {
    local interface ObjectAdapter {
        string getName();
        Communicator getCommunicator();

        // ...
    };

    local interface Communicator {
        ObjectAdapter createObjectAdapter(string name);
        ObjectAdapter createObjectAdapterWithEndpoints(string name, string
endpoints);
        ObjectAdapter createObjectAdapterWithRouter(string name, Router* rtr);

        // ...
    };
};
```

The `ObjectAdapter` operations behave as follows:

- The `getName` operation returns the name of the adapter as passed to one of the communicator operations `createObjectAdapter`, `createObjectAdapterWithEndpoints`, or `createObjectAdapterWithRouter`.
- The `getCommunicator` operation returns the communicator that was used to create the adapter.

Note that there are other operations in the `ObjectAdapter` interface; we will explore these throughout the remainder of this discussion.

An application normally needs to configure an object adapter with endpoints or a router. Calling `createObjectAdapter` with a non-empty value for `name` means the new object adapter will check the communicator's configuration for properties, including:

- *`name`*`.Endpoints` - defines one or more object adapter endpoints
- *`name`*`.Router` - specifies the stringified proxy of a router

`createObjectAdapter` raises an exception if neither is defined. Passing an empty value for `name` to `createObjectAdapter` creates an object adapter that can only be used for collocated invocations.

The `createObjectAdapterWithEndpoints` operation allows you to specify the object adapter's endpoints directly, overriding any setting for *`name`*`.Endpoints`. Similarly, `createObjectAdapterWithRouter` accepts a router proxy and overrides *`name`*`.Router`.

See Also

- Communicators

Copyright 2016, ZeroC, Inc.

Servant Activation and Deactivation

The term *servant activation* refers to making the presence of a servant for a particular Ice object known to the Ice run time. Activating a servant adds an entry to the Active Servant Map (ASM). Another way of looking at servant activation is to think of it as creating a link between the identity of an Ice object and the corresponding programming-language servant that handles requests for that Ice object. Once the Ice run time has knowledge of this link, it can dispatch incoming requests to the correct servant. Without this link, that is, without a corresponding entry in the ASM, an incoming request for the identity results in an `ObjectNotExistException`. While a servant is activated, it is said to *incarnate* the corresponding Ice object.

The inverse operation is known as *servant deactivation*. Deactivating a servant removes an entry for a particular identity from the ASM. Thereafter, incoming requests for that identity are no longer dispatched to the servant and result in an `ObjectNotExistException`.

The object adapter offers a number of operations for managing servant activation and deactivation:

Slice

```
module Ice {
    local interface ObjectAdapter {
        // ...

        Object* add(Object servant, Identity id);
        Object* addWithUUID(Object servant);
        Object  remove(Identity id);
        Object  find(Identity id);
        Object  findByProxy(Object* proxy);

        // ...
    };
};
```

The operations behave as follows:

- `add`

  The `add` operation adds a servant with the given identity to the ASM. Requests are dispatched to that servant as soon as `add` is called. The return value is the proxy for the Ice object incarnated by that servant. The proxy embeds the identity passed to `add`.

  You cannot call `add` with the same identity more than once: attempts to add an already existing identity to the ASM result in an `AlreadyRegisteredException`. (It does not make sense to add two servants with the same identity because that would make it ambiguous as to which servant should handle incoming requests for that identity.)

  Note that it is possible to activate the same servant multiple times with different identities. In that case, the same single servant incarnates multiple Ice objects. We explore the ramifications of this in more detail in our discussion of server implementation techniques.

- `addWithUUID`

  The `addWithUUID` operation behaves the same way as the `add` operation but does not require you to supply an identity for the servant. Instead, `addWithUUID` generates a UUID as the identity for the corresponding Ice object. You can retrieve the generated identity by calling the `ice_getIdentity` operation on the returned proxy. `addWithUUID` is useful to create identities for temporary objects, such as short-lived session objects. (You can also use `addWithUUID` for persistent objects that do not have a natural identity, as we have done for the file system application.)

- `remove`

  The `remove` operation breaks the association between an identity and its servant by removing the corresponding entry from the ASM; it returns a reference to the removed servant.

  Once the servant is deactivated, new incoming requests for the removed identity cause the client to receive an `ObjectNotExistException`. Requests that are executing inside the servant at the time `remove` is called are allowed to complete normally. Once the

Copyright 2016, ZeroC, Inc.

last request for the servant is complete, the object adapter drops its reference (or smart pointer, for C++) to the servant. At that point, the servant becomes available for garbage collection (or is destroyed, for C++), provided there are no other references or smart pointers to the servant. The net effect is that a deactivated servant is destroyed once it becomes idle.

Deactivating an object adapter implicitly calls `remove` on its active servants.

- `find`

  The `find` operation performs a lookup in the ASM and returns the servant for the specified object identity. If no servant with that identity is registered, the operation returns null. Note that `find` does not consult any servant locators or default servants.

- `findByProxy`

  The `findByProxy` operation performs a lookup in the ASM and returns the servant with the object identity and facet that are embedded in the proxy. If no such servant is registered, the operation returns null. Note that `findByProxy` does not consult any servant locators or default servants.

See Also

- The Active Servant Map
- Object Identity
- Object Adapter States
- Server Implementation Techniques
- Servant Locators
- Default Servants

Copyright 2016, ZeroC, Inc.

Object Adapter States

On this page:

- Object Adapter State Transitions
- Changing Object Adapter States

Object Adapter State Transitions

An object adapter has a number of processing states:

- **Holding**

  In this state, any incoming requests for the adapter are held, that is, not dispatched to servants.

  For TCP/IP (and other stream-oriented protocols), the server-side run time stops reading from the corresponding transport endpoint while the adapter is in the holding state. In addition, it also does not accept incoming connection requests from clients. This means that if a client sends a request to an adapter that is in the holding state, the client eventually receives a `TimeoutException` or `ConnectTimeoutException` (unless the adapter is placed into the active state before the timer expires).

  For UDP, client requests that arrive at an adapter that is in the holding state are thrown away.

  Immediately after creation of an adapter, the adapter is in the holding state. This means that requests are not dispatched until you place the adapter into the active state.

  Note that bidirectional adapters cannot be placed into the holding state. If you call `hold` on a bidirectional adapter, the call does nothing.

- **Active**

  In this state, the adapter accepts incoming requests and dispatches them to servants. A newly-created adapter is initially in the holding state. The adapter begins dispatching requests as soon as you place it into the active state.

  You can transition between the active and the holding state as many times as you wish.

  Note that bidirectional adapters need not be activated. Further, calls to collocated servants (that is, to servants that are activated in the communicator that created the proxy) succeed even if the adapter is not activated, unless you have disabled collocation optimization.

- **Inactive**

  In this state, the adapter has conceptually been destroyed (or is in the process of being destroyed). Deactivating an adapter destroys all transport endpoints that are associated with the adapter. Requests that are executing at the time the adapter is placed into the inactive state are allowed to complete, but no new requests are accepted. (New requests are rejected with an exception). Any attempt to use a deactivated object adapter results in an `ObjectAdapterDeactivatedException`.

Changing Object Adapter States

The `ObjectAdapter` interface offers operations that allow you to change the adapter state, as well as to wait for a state change to be complete:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local interface ObjectAdapter {
        // ...

        void activate();
        void hold();
        void waitForHold();
        void deactivate();
        void waitForDeactivate();
        void isDeactivated();
        void destroy();

        // ...
    };
};
```

The operations behave as follows:

- `activate`

  The `activate` operation places the adapter into the active state. Activating an adapter that is already active has no effect. The Ice run time starts dispatching requests to servants for the adapter as soon as `activate` is called.

- `hold`

  The `hold` operation places the adapter into the holding state. Requests that arrive after calling `hold` are held as described above. Requests that are in progress at the time `hold` is called are allowed to complete normally. Note that `hold` returns immediately without waiting for currently executing requests to complete.

- `waitForHold`

  The `waitForHold` operation suspends the calling thread until the adapter has completed its transition to the holding state, that is, until all currently executing requests have finished. You can call `waitForHold` from multiple threads, and you can call `waitForHold` while the adapter is in the active state. If you call `waitForHold` on an adapter that is already in the holding state, `waitForHold` returns immediately.

- `deactivate`

  The `deactivate` operation initiates deactivation of the adapter: requests that arrive after calling `deactivate` are rejected, but currently executing requests are allowed to complete. Once all requests have completed, the adapter enters the inactive state. Note that `deactivate` returns immediately without waiting for the currently executing requests to complete. A deactivated adapter cannot be reactivated; you can create a new adapter with the same name, but only after calling `destroy` on the existing adapter. Any attempt to use a deactivated object adapter results in an `ObjectAdapterDeactivatedException`.

- `waitForDeactivate`

  The `waitForDeactivate` operation suspends the calling thread until the adapter has completed its transition to the inactive state, that is, until all currently executing requests have completed. You can call `waitForDeactivate` from multiple threads, and you can call `waitForDeactivate` while the adapter is in the active or holding state. Calling `waitForDeactivate` on an adapter that is in the inactive state does nothing and returns immediately.

- `isDeactivated`

  The `isDeactivated` operation returns true if `deactivate` has been invoked on the adapter. A return value of true does not necessarily indicate that the adapter has fully transitioned to the inactive state, only that it has begun this transition. Applications that need to know when deactivation is completed can use `waitForDeactivate`.

Copyright 2016, ZeroC, Inc.

- `destroy`

  The `destroy` operation deactivates the adapter and releases all of its resources. Internally, `destroy` invokes `deactivate` followe d by `waitForDeactivate`, therefore the operation blocks until all currently executing requests have completed. Furthermore, any servants associated with the adapter are destroyed, all transport endpoints are closed, and the adapter's name becomes available for reuse.

  Destroying a communicator implicitly destroys all of its object adapters. Invoking `destroy` on an adapter is only necessary when you need to ensure that its resources are released prior to the destruction of its communicator.

Placing an adapter into the holding state is useful, for example, if you need to make state changes in the server that require the server (or a group of servants) to be idle. For example, you could place the implementation of your servants into a dynamic library and upgrade the implementation by loading a newer version of the library at run time without having to shut down the server.

Similarly, waiting for an adapter to complete its transition to the inactive state is useful if your server needs to perform some final clean-up work that cannot be carried out until all executing requests have completed.

Note that you can create an object adapter with the same name as a previous object adapter, but only once `destroy` on the previous adapter has completed.

See Also

- Collocated Invocation and Dispatch
- Bidirectional Connections

Copyright 2016, ZeroC, Inc.

Object Adapter Endpoints

An object adapter maintains two sets of transport endpoints. One set identifies the network interfaces on which the adapter listens for new connections, and the other set is embedded in proxies created by the adapter and used by clients to communicate with it. We will refer to these sets of endpoints as the *physical endpoints* and the *published endpoints*, respectively. In most cases these sets are identical, but there are situations when they must be configured independently.

On this page:

- Physical Object Adapter Endpoints
- Published Object Adapter Endpoints
- Refreshing Object Adapter Endpoints
- Timeouts in Object Adapter Endpoints
- Discovering Object Adapter Endpoints
- A Router's Effect on Object Adapter Endpoints

Physical Object Adapter Endpoints

An object adapter's physical endpoints identify the network interfaces on which it receives requests from clients. These endpoints are configured via the *name*`.Endpoints` property, or they can be specified explicitly when creating an adapter using the operation `createObjectAdapterWithEndpoints`. The endpoint syntax generally consists of a transport protocol followed by an optional host name and port.

If a host name is specified, the object adapter listens only on the network interface associated with that host name. If no host name is specified but the property `Ice.Default.Host` is defined, the object adapter uses the property's value as the host name. Finally, if a host name is not specified, and the property `Ice.Default.Host` is undefined, the object adapter listens on all available network interfaces, including the loopback interface. You may also force the object adapter to listen on all interfaces by using one of the host names `0.0.0.0` or `*`. The adapter does *not* expand the list of interfaces when it is initialized. Instead, if no host is specified, or you use `-h *` or `-h 0.0.0.0`, the adapter binds to `INADDR_ANY` to listen for incoming requests.

If you want an adapter to accept requests on certain network interfaces, you must specify a separate endpoint for each interface. For example, the following property configures a single endpoint for the adapter named `MyAdapter`:

```
MyAdapter.Endpoints=tcp -h 10.0.1.1 -p 9999
```

This endpoint causes the adapter to accept requests on the network interface associated with the IP address `10.0.1.1` at port `9999`. Note however that this adapter configuration does not accept requests on the loopback interface (the one associated with address `127.0.0.1`). If both addresses must be supported, then both must be specified explicitly, as shown below:

```
MyAdapter.Endpoints=tcp -h 10.0.1.1 -p 9999:tcp -h 127.0.0.1 -p 9999
```

If these are the only two network interfaces available on the host, then a simpler configuration omits the host name altogether, causing the object adapter to listen on both interfaces automatically:

```
MyAdapter.Endpoints=tcp -p 9999
```

If you want to make your configuration more explicit, you can use one of the special host names mentioned earlier:

```
MyAdapter.Endpoints=tcp -h * -p 9999
```

One advantage of this last example is that it ensures the object adapter *always* listens on all interfaces, even if a definition for `Ice.Default.Host` is later added to your configuration. Without an explicit host name, the addition of `Ice.Default.Host` could potentially change the

Copyright 2016, ZeroC, Inc.

interfaces on which the adapter is listening. For diagnostic purposes, you can determine the set of local addresses that Ice substitutes for the wildcard address by setting the property `Ice.Trace.Network=3` and reviewing the server's log output.

Careful consideration must also be given to the selection of a port for an endpoint. If no port is specified, the adapter uses a port that is selected (essentially at random) by the operating system, meaning the adapter will likely be using a different port each time the server is restarted. Whether that behavior is desirable depends on the application, but in many applications a client has a proxy containing the adapter's endpoint and expects that proxy to remain valid indefinitely. Therefore, an endpoint generally should contain a fixed port to ensure that the adapter is always listening at the same port.

However, there are certain situations where a fixed port is not required. For example, an adapter whose servants are transient does not need a fixed port, because the proxies for those objects are not expected to remain valid past the lifetime of the server process. Similarly, a server using indirect binding via IceGrid does not need a fixed port because its port is never published.

Published Object Adapter Endpoints

An object adapter publishes its endpoints in the proxies it creates, but it is not always appropriate to publish the adapter's physical endpoints in a proxy. For example, suppose a server is running on a host in a private network, protected from the public network by a firewall that can forward network traffic to the server. The adapter's physical endpoints must use the private network's address scheme, but a client in the public network would be unable to use those endpoints if they were published in a proxy. In this scenario, the adapter must publish endpoints in its proxies that direct the client to the firewall instead.

The published endpoints are configured using the adapter property *name*`.PublishedEndpoints`. If this property is not defined, the adapter publishes its physical endpoints by default, with one exception: endpoints for the loopback address (`127.0.0.1`) are not published unless the loopback interface is the only interface, or `127.0.0.1` (or `loopback`) is explicitly listed as an endpoint with the `-h` option. Otherwise, to force the inclusion of loopback endpoints when they would normally be excluded, you must define *name*`.PublishedEndpoints` explicitly.

As an example, the properties below configure the adapter named `MyAdapter` with physical and published endpoints:

```
MyAdapter.Endpoints=tcp -h 10.0.1.1 -p 9999
MyAdapter.PublishedEndpoints=tcp -h corpfw -p 25000
```

This example assumes that clients connecting to host `corpfw` at port `25000` are forwarded to the adapter's endpoint in the private network.

Another use case of published endpoints is for replicated servers. Suppose we have two instances of a stateless server running on separate hosts in order to distribute the load between them. We can supply the client with a bootstrap proxy containing the endpoints of both servers, and the Ice run time in the client will select one of the servers at random when a connection is established. However, should the client invoke an operation on a server that returns a proxy for another object, that proxy would normally contain only the endpoint of the server that created it. Invocations on the new proxy would always be directed at the same server, reducing the opportunity for load balancing.

We can alleviate this situation by configuring the adapters to publish the endpoints of both servers. For example, here is a configuration for the server on host `Sun1`:

```
MyAdapter.Endpoints=tcp -h Sun1 -p 9999
MyAdapter.PublishedEndpoints=tcp -h Sun1 -p 9999:tcp -h Sun2 -p 9999
```

Similarly, the configuration for host `Sun2` retains the same published endpoints:

```
MyAdapter.Endpoints=tcp -h Sun2 -p 9999
MyAdapter.PublishedEndpoints=tcp -h Sun1 -p 9999:tcp -h Sun2 -p 9999
```

For troubleshooting purposes, you can examine the published endpoints for an object adapter by setting the property `Ice.Trace.Network=3`. Note however that this setting generates significant trace information about the Ice run time's network activity, therefore you may not want to use this setting by default.

Copyright 2016, ZeroC, Inc.

Refreshing Object Adapter Endpoints

A host's network interfaces may change over time, for example, if a laptop moves in and out of range of a wireless network. The object adapter provides an operation to refresh its list of interfaces:

Slice

```
local interface ObjectAdapter {
    void refreshPublishedEndpoints();
    // ...
};
```

Calling `refreshPublishedEndpoints` causes the object adapter to update its internal list of available network interfaces and to use this updated information in the *`name`*`.PublishedEndpoints` property. This allows you to react to changing network interfaces while an object adapter is in use, but your application code is responsible for determining when it is necessary to call this operation.

Note that `refreshPublishedEndpoints` takes effect only for object adapters that specify published endpoints without a host, or that set the published endpoints to `-h *` or `-h 0.0.0.0`.

Timeouts in Object Adapter Endpoints

As a defense against hostile clients, we recommend that you specify a timeout for your physical object adapter endpoints. The timeout value you select affects tasks that the Ice run time normally does not expect to block for any significant amount of time, such as writing a reply message to a socket or waiting for SSL negotiation to complete. If you don't specify a timeout for an endpoint, the Ice run time uses the value of the `Ice.Default.Timeout` property. You can also specify the value `infinite` in an endpoint to wait indefinitely, but be aware that this could allow malicious or misbehaving clients to consume excessive resources such as file descriptors.

Specifying a timeout in an object adapter endpoint is done exactly as in a proxy endpoint, using the `-t` option:

```
MyAdapter.Endpoints=tcp -p 9999 -t 5000
```

In this example, we specify a timeout of five seconds.

Unless overridden by published endpoints, the timeout specified in your object adapter endpoint also appears in the endpoints of any proxies created by that object adapter. This feature allows you to provide a default timeout value for clients that use your object adapter, although the client's code or configuration can cause it to use a different timeout in practice.

Discovering Object Adapter Endpoints

The object adapter provides operations for retrieving the physical and published endpoints:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local interface ObjectAdapter {
        // ...

        EndpointSeq getEndpoints();
        EndpointSeq getPublishedEndpoints();

        // ...
    };
};
```

The sequences that are returned contain `Endpoint` objects representing the adapter's physical and published endpoints, respectively.

A Router's Effect on Object Adapter Endpoints

If an object adapter is configured with a router, the adapter's published endpoints are augmented to reflect the router. See Glacier2 for more information on configuring an adapter with a router.

See Also

- Proxy and Endpoint Syntax
- Object Adapter Properties
- Ice.Default.*
- IceGrid
- Using Connections
- Communicators
- Glacier2

Copyright 2016, ZeroC, Inc.

Creating Proxies with an Object Adapter

Proxies are created as a side-effect of using the servant activation operations, but the life cycle of proxies is completely independent from that of servants. The `ObjectAdapter` interface provides several operations for creating a proxy for an object, regardless of whether a servant is currently activated for that object's identity:

Slice

```
module Ice {
    local interface ObjectAdapter {
        // ...

        Object* createProxy(Identity id);
        Object* createDirectProxy(Identity id);
        Object* createIndirectProxy(Identity id);

        // ...
    };
};
```

These operations are described below:

- `createProxy`

  The `createProxy` operation returns a new proxy for the object with the given identity. The adapter's configuration determines whether the return value is a direct proxy or an indirect proxy. If the adapter is configured with an adapter ID, the operation returns an indirect proxy that refers to the adapter ID. If the adapter is also configured with a replica group ID, the operation returns an indirect proxy that refers to the replica group ID. Otherwise, if an adapter ID is not defined, `createProxy` returns a direct proxy containing the adapter's published endpoints.

- `createDirectProxy`

  The `createDirectProxy` operation returns a direct proxy containing the adapter's published endpoints.

- `createIndirectProxy`

  The `createIndirectProxy` operation returns an indirect proxy. If the adapter is configured with an adapter ID, the returned proxy refers to that adapter ID. Otherwise, the proxy refers only to the object's identity.

In contrast to `createProxy`, `createIndirectProxy` does not use the replica group ID. Therefore, the returned proxy always refers to a specific replica.

After using one of the operations discussed above to create a proxy, you will receive a proxy that is configured by default for twoway invocations. If you require the proxy to have a different configuration, you can use the proxy factory methods to create a new proxy with the desired configuration. As an example, the code below demonstrates how to configure the proxy for oneway invocations:

C++

```
Ice::ObjectAdapterPtr adapter = ...;
Ice::Identity id = ...;
Ice::ObjectPrx proxy = adapter->createProxy(id)->ice_oneway();
```

You can also instruct the object adapter to use a different default proxy configuration by setting the property *`name`*`.ProxyOptions`. For example, the following property causes the object adapter to return proxies that are configured for oneway invocations by default:

Copyright 2016, ZeroC, Inc.

```
MyAdapter.ProxyOptions=-o
```

See Also

- Servant Activation and Deactivation
- Object Life Cycle
- Object Identity
- Terminology
- Object Adapter Endpoints
- Object Adapter Replication
- Proxy Methods
- Object Adapter Properties

Copyright 2016, ZeroC, Inc.

Using Multiple Object Adapters

A typical server rarely needs to use more than one object adapter. If you are considering using multiple object adapters, we suggest that you check whether any of the considerations in the list below apply to your situation:

- You need fine-grained control over which objects are accessible. For example, you could have an object adapter with only secure endpoints to restrict access to some administrative objects, and another object adapter with non-secure endpoints for other objects. Because an object adapter is associated with one or more transport endpoints, you can firewall a particular port, so objects associated with the corresponding endpoint cannot be reached unless the firewall rules are satisfied.
- You need control over the number of threads in the pools for different sets of objects in your application. For example, you may not need concurrency on the objects connected to a particular object adapter, and multiple object adapters, each with its own thread pool, can be useful to solve deadlocks.
- You want to be able to temporarily disable processing new requests for a set of objects. This can be accomplished by placing an object adapter in the holding state.
- You want to set up different request routing when using an Ice router with Glacier2.

If none of the preceding items apply, chances are that you do not need more than one object adapter.

See Also

- Object Adapter Thread Pools
- Nested Invocations
- Object Adapter States
- Glacier2

Copyright 2016, ZeroC, Inc.

The Current Object

Up to now, we have tacitly ignored the trailing parameter of type `Ice::Current` that is passed to each skeleton operation on the server side. The `Current` structure is defined as follows:

Slice

```
module Ice {
    local dictionary<string, string> Context;

    enum OperationMode { Normal, \Idempotent };

    local struct Current {
        ObjectAdapter   adapter;
        Connection      con;
        Identity        id;
        string          facet;
        string          operation;
        OperationMode   mode;
        Context         ctx;
        int             requestId;
        EncodingVersion encoding;
    };
};
```

Note that the `Current` value provides access to information about the currently executing request to the implementation of an operation in the server:

- `adapter`
  The `adapter` member provides access to the object adapter via which the request is being dispatched. In turn, the adapter provides access to its communicator (via the `getCommunicator` operation).
- `con`
  The `con` member provides information about the connection over which this request was received.
- `id`
  The `id` member provides the object identity for the current request.
- `facet`
  The `facet` member provides access to the facet for the request.
- `operation`
  The `operation` member contains the name of the operation that is being invoked. Note that the operation name may indicate one of the operations on `Ice::Object`, such as `ice_ping` or `ice_isA`. (`ice_isA` is invoked if a client performs a `checkedCast`.)
- `mode`
  The `mode` member contains the invocation mode for the operation (`Normal` or `Idempotent`), which influences the retry behavior of the Ice run time.
- `ctx`
  The `ctx` member contains the current request context for the invocation.
- `requestId`
  The Ice protocol uses request IDs to associate replies with their corresponding requests. The `requestId` member provides this ID. For oneway requests (which do not have replies), the request ID is `0`. For collocated requests (which do not use the Ice protocol), the request ID is `-1`.
- `encoding`
  The encoding version used to encode the input and output parameters.

If you implement your server such that it uses a separate servant for each Ice object, the contents of `Current` are not particularly

Copyright 2016, ZeroC, Inc.

interesting. (You would most likely access `Current` to read the `adapter` member, for example, to activate or deactivate a servant.) However, as we will see in our discussion of default servants and servant locators, the `Current` object is essential for more sophisticated (and more scalable) servant implementations.

See Also

- Object Identity
- Default Servants
- Servant Locators
- Request Contexts
- Versioning
- Using Connections
- The Ice Protocol

Copyright 2016, ZeroC, Inc.

Servant Locators

In a nutshell, a servant locator is a local object that you implement and attach to an object adapter. Once an adapter has a servant locator, it consults its active servant map (ASM) to locate a servant for an incoming request as usual. If a servant for the request can be found in the ASM, the request is dispatched to that servant. However, if the ASM does not have an entry for the object identity of the request, the object adapter calls back into the servant locator to ask it to provide a servant for the request. The servant locator either:

- locates an existing servant or instantiates a new servant and passes it to the Ice run time, in which case the request is dispatched to that servant, or
- the servant locator indicates failure to locate a servant to the Ice run time, in which case the client receives an `ObjectNotExistException`.

This simple mechanism allows us to scale servers to provide access to an unlimited number of Ice objects. For example, instead of instantiating a separate servant for each and every Ice object in existence, the server can instantiate servants only for those Ice objects that are actually used by clients. Similar to a Default Servant, the server can also instantiate only one single servant that the locator returns for a group of Ice objects with similar properties.

Servant locators are most commonly used by servers that provide access to databases: typically, the number of entries in the database is far larger than what the server can hold in memory. Servant locators allow the server to only instantiate servants for those Ice objects that are actually used by clients.

Another common use for servant locators is in servers that are used for process control or network management: in that case, there is no database but, instead, there is a potentially very large number of devices or network elements that must be controlled via the server. Otherwise, this scenario is the same as for large databases: the number of Ice objects exceeds the number of servants that the server can hold in memory and, therefore, requires an approach that allows the number of instantiated servants to be less than the number of Ice objects.

Topics

- The ServantLocator Interface
- Threading Guarantees for Servant Locators
- Registering a Servant Locator
- Servant Locator Example
- Using Identity Categories with Servant Locators
- Using Cookies with Servant Locators

See Also

- Object Adapters
- The Active Servant Map
- Object Identity
- Object Life Cycle
- Default Servants

Copyright 2016, ZeroC, Inc.

The ServantLocator Interface

A servant locator has the following interface:

Slice

```
module Ice {
    local interface ServantLocator {
        ["UserException"]
        Object locate(Current curr, out LocalObject cookie);

        ["UserException"]
        void finished(Current curr, Object servant, LocalObject cookie);

        void deactivate(string category);
    };
};
```

Note that `ServantLocator` is a local interface. To create an actual implementation of a servant locator, you must define a class that is derived from `Ice::ServantLocator` and provide implementations of the `locate`, `finished`, and `deactivate` operations. The Ice run time invokes the operations on your derived class as follows:

- `locate`

  Whenever a request arrives for which no entry exists in the active servant map (ASM), the Ice run time calls `locate` and supplies the `Current` object for the request. The implementation of `locate` (which you provide as part of the derived class) is supposed to return a servant that can process the incoming request. Your implementation of `locate` can behave in three possible ways:

  1. Locate an existing or instantiate a new servant, and return this servant for the current request. In this case, the Ice run time dispatches the request to this servant.
  2. `ObjectNotExistException`
  3. `ObjectNotExistExceptionOperationNotExistExceptionFacetNotExistExceptionUnknownLocalExceptio n`

  You can also throw user exceptions from `locate`. If the user exception is in the corresponding operation's exception specification, that user exception is returned to the client. User exceptions thrown by `locate` that are not listed in the exception specification of the corresponding operation are also returned to the client, and then thrown in the client as `UnknownUserException`. Non-Ice exceptions are returned to the client as `UnknownException`.

  The `cookie` out-parameter to `locate` allows you to return a local object to the object adapter. The object adapter does not care about the contents of that object (and it is legal to return a null cookie). Instead, the Ice run time passes whatever cookie you return from `locate` back to you when it calls `finished`. This allows you to pass an arbitrary amount of state from `locate` to the corresponding call to `finished`.

- `finished`

  If a call to `locate` has returned a servant to the Ice run time, the Ice run time dispatches the incoming request to the servant. Once the request is complete (that is, the operation being invoked has completed), the Ice run time calls `finished`, passing the servant whose operation has completed, the `Current` object for the request, and the cookie that was initially created by `locate`. This means that every call to `locate` is balanced by a corresponding call to `finished` (provided that `locate` actually returned a servant).

  If you throw an exception from `finished`, the Ice run time propagates the thrown exception back to the client. As for `locate`, you can throw user exceptions from `finished`. If a user exception is in the corresponding operation's exception specification, that user exception is returned to the client. User exceptions that are not in the corresponding operation's exception specification are also returned to the client, and then thrown by the client as `UnknownUserException`.

  `finished` can also throw run-time exceptions. However, only `ObjectNotExistException`, `OperationNotExistException`, and `FacetNotExistException` are propagated without change to the client; other run-time exceptions are returned to the client as `UnknownLocalException`.

Copyright 2016, ZeroC, Inc.

Non-Ice exceptions thrown from `finished` are returned to the client as `UnknownException`.

If both the operation implementation and `finished` throw a user exception, the exception thrown by `finished` overrides the exception thrown by the operation.

- `deactivate`

  The `deactivate` operation allows a servant locator to clean up once it is no longer needed. (For example, the locator might close a database connection.) The Ice run time passes the category of the servant locator being deactivated to the `deactivate` operation.

  The run time calls `deactivate` when destroying the object adapter to which the servant locator is attached. More precisely, `deactivate` is called when you call `destroy` on the object adapter, or when you call `destroy` on the communicator (which implicitly calls `destroy` on the object adapter).

  Once the run time has called `deactivate`, it is guaranteed that no further calls to `locate` or `finished` can happen, that is, `deactivate` is called exactly once, after all operations dispatched via this servant locator have completed.

  This also explains why `deactivate` is not called as part of `ObjectAdapter::deactivate`: `ObjectAdapter::deactivate` initiates deactivation and returns immediately, so it cannot call `ServantLocator::deactivate` directly, because there might still be outstanding requests dispatched via this servant locator that have to complete first — in turn, this would mean that either `ObjectAdapter::deactivate` could block (which it must not do) or that a call to `ServantLocator::deactivate` could be followed by one or more calls to `finished` (which must not happen either).

It is important to realize that the Ice run time does not "remember" the servant that is returned by a particular call to `locate`. Instead, the Ice run time simply dispatches an incoming request to the servant returned by `locate` and, once the request is complete, calls `finished`. In particular, if two requests for the same servant arrive more or less simultaneously, the Ice run time calls `locate` and `finished` once for each request. In other words, `locate` establishes the association between an object identity and a servant; that association is valid only for a single request and is never used by the Ice run time to dispatch a different request.

See Also

- The Active Servant Map
- The Current Object
- Run-Time Exceptions
- Object Adapter States

Copyright 2016, ZeroC, Inc.

Threading Guarantees for Servant Locators

The Ice run time guarantees that every operation invocation that involves a servant locator is bracketed by calls to `locate` and `finished`, that is, every call to `locate` is balanced by a corresponding call to `finished` (assuming that the call to `locate` actually returned a servant, of course).

In addition, the Ice run time guarantees that `locate`, the operation, and `finished` are called by the same thread. This guarantee is important because it allows you to use `locate` and `finished` to implement thread-specific pre- and post-processing around operation invocations. (For example, you can start a transaction in `locate` and commit or roll back that transaction in `finished`, or you can acquire a lock in `locate` and release the lock in `finished`.

Both transactions and locks usually are thread-specific, that is, only the thread that started a transaction can commit it or roll it back, and only the thread that acquired a lock can release the lock.

If you are using asynchronous method dispatch, the thread that starts a call is not necessarily the thread that finishes it. In that case, `finished` is called by whatever thread executes the operation implementation, which may be a different thread than the one that called `locate`.

The Ice run time also guarantees that `deactivate` is called when you destroy the object adapter to which the servant locator is attached. The `deactivate` call is made only once all operations that involved the servant locator are finished, that is, `deactivate` is guaranteed not to run concurrently with `locate` or `finished`, and is guaranteed to be the last call made to a servant locator.

Beyond this, the Ice run time provides no threading guarantees for servant locators. In particular, it is possible for invocations of:

- `locate` to proceed concurrently (for the same object identity or for different object identities).
- `finished` to proceed concurrently (for the same object identity or for different object identities).
- `locate` and `finished` to proceed concurrently (for the same object identity or for different object identities).

These semantics allow you to extract the maximum amount of parallelism from your application code (because the Ice run time does not serialize invocations when serialization may not be necessary). Of course, this means that you must protect access to shared data from `locate` and `finished` with mutual exclusion primitives as necessary.

See Also

- The ServantLocator Interface
- The Ice Threading Model

Copyright 2016, ZeroC, Inc.

Registering a Servant Locator

On this page:

- Servant Locator Registration
- Call Dispatch Semantics for Servant Locators

Servant Locator Registration

An object adapter does not automatically know when you create a servant locator. Instead, you must explicitly register a servant locator with the object adapter:

Slice

```
module Ice {
    local interface ObjectAdapter {
        // ...

        void addServantLocator(ServantLocator locator, string category);

        ServantLocator removeServantLocator(string category);

        ServantLocator findServantLocator(string category);

        // ...
    };
};
```

As you can see, the object adapter allows you to add, remove, and find servant locators. Note that, when you register a servant locator, you must provide an argument for the `category` parameter. The value of the `category` parameter controls which object identities the servant locator is responsible for: only object identities with a matching `category` member trigger a corresponding call to `locate`. An incoming request for which no explicit entry exists in the active servant map (ASM) and with a category for which no servant locator is registered returns an `ObjectNotExistException` to the client.

`addServantLocator` has the following semantics:

- You can register exactly one servant locator for a specific category. Attempts to call `addServantLocator` for the same category more than once raise an `AlreadyRegisteredException`.
- You can register different servant locators for different categories, or you can register the same single servant locator multiple times (each time for a different category). In the former case, the category is implicit in the servant locator instance that is called by the Ice run time; in the latter case, the implementation of `locate` can find out which category the incoming request is for by examining the object identity member of the `Current` object that is passed to `locate`.
- It is legal to register a servant locator for the empty category. Such a servant locator is known as a *default servant locator*: if a request comes in for which no entry exists in the ASM, and whose category does not match the category of any other registered servant locator, the Ice run time calls `locate` on the default servant locator.

`removeServantLocator` removes and returns the servant locator for a specific category (including the empty category) with the following semantics:

- If no servant locator is registered for the specified category, the operation raises `NotRegisteredException`.
- Once a servant locator is successfully removed for the specified category, the Ice run time guarantees that no new incoming requests for that category are dispatched to the servant locator.
- A call to `removeServantLocator` returns immediately without waiting for the completion of any pending requests on that servant locator; such requests still complete normally by calling `finished` on the servant locator.
- Removing a servant locator does not cause Ice to invoke `deactivate` on that servant locator, as `deactivate` is only called when a registered servant locator's object adapter is destroyed.

`findServantLocator` allows you to retrieve the servant locator for a specific category (including the empty category). If no match is found, the operation returns null.

Copyright 2016, ZeroC, Inc.

Call Dispatch Semantics for Servant Locators

The preceding rules may seem complicated, so here is a summary of the actions taken by the Ice run time to locate a servant for an incoming request.

Every incoming request implicitly identifies a specific object adapter for the request (because the request arrives at a specific transport endpoint and, therefore, identifies a particular object adapter). The incoming request carries an object identity that must be mapped to a servant. To locate a servant, the Ice run time goes through the following steps, in the order shown:

1. Look for the identity in the ASM. If the ASM contains an entry, dispatch the request to the corresponding servant.
2. If the category of the incoming object identity is non-empty, look for a default servant that is registered for that category. If such a default servant is registered, dispatch the request to that servant.
3. If the category of the incoming object identity is empty, or no default servant could be found for the category in step 2, look for a default servant that is registered for the empty category. If such a default servant is registered, dispatch the request to that servant.
4. If the category of the incoming object identity is non-empty and no servant could be found in the preceding steps, look for a servant locator that is registered for that category. If such a servant locator is registered, call `locate` on the servant locator and, if `locate` returns a servant, dispatch the request to that servant, followed by a call to `finished`; otherwise, if the call to `locate` returns null, raise `ObjectNotExistException` or `FacetNotExistException` in the client.
5. If the category of the incoming object identity is empty, or no servant locator could be found for the category in step 4, look for a default servant locator (that is, a servant locator that is registered for the empty category). If a default servant locator is registered, dispatch the request as for step 4.
6. Raise `ObjectNotExistException` or `FacetNotExistException` in the client. (`ObjectNotExistException` is raised if the ASM does not contain a servant with the given identity at all, `FacetNotExistException` is raised if the ASM contains a servant with a matching identity, but a non-matching facet.)

It is important to keep these call dispatch semantics in mind because they enable a number of powerful implementation techniques. Each technique allows you to streamline your server implementation and to precisely control the trade-off between performance, memory consumption, and scalability. To illustrate the possibilities, we will outline a number of the most common implementation techniques.

See Also

- Object Adapters
- Object Identity
- The Active Servant Map
- Default Servants
- Versioning
- Servant Locator Example

Copyright 2016, ZeroC, Inc.

Servant Locator Example

To illustrate the servant locator concepts outlined so far, let us examine a (very simple) implementation. Consider that we want to create an electronic phone book for the entire world's telephone system (which, clearly, involves a very large number of entries, certainly too many to hold the entire phone book in memory). The actual phone book entries are kept in a large database. Also assume that we have a search operation that returns the details of a phone book entry. The Slice definitions for this application might look something like the following:

Slice

```
struct Details {
    // Lots of details about the entry here...
};

interface PhoneEntry {
    idempotent Details getDetails();
    idempotent void updateDetails(Details d);
    // ...
};

struct SearchCriteria {
    // Fields to permit searching...
};

interface PhoneBook {
    idempotent PhoneEntry* search(SearchCriteria c);
    // ...
};
```

The details of the application do not really matter here; the important point to note is that each phone book entry is represented as an interface for which we need to create a servant eventually, but we cannot afford to keep servants for all entries permanently in memory.

Each entry in the phone database has a unique identifier. This identifier might be an internal database identifier, or a combination of field values, depending on exactly how the database is constructed. The important point is that we can use this database identifier to link the proxy for an Ice object to its persistent state: we simply use the database identifier as the object identity. This means that each proxy contains the primary access key that is required to locate the persistent state of each Ice object and, therefore, instantiate a servant for that Ice object.

What follows is an outline implementation in C++. The class definition of our servant locator looks as follows:

C++

```
class MyServantLocator : public virtual Ice::ServantLocator {
public:

    virtual Ice::ObjectPtr locate(const Ice::Current& c,
Ice::LocalObjectPtr& cookie);

    virtual void finished(const Ice::Current& c, const Ice::ObjectPtr& servant,
                          const Ice::LocalObjectPtr& cookie);

    virtual void deactivate(const std::string& category);
};
```

Note that `MyServantLocator` inherits from `Ice::ServantLocator` and implements the pure virtual functions that are generated by the s

Copyright 2016, ZeroC, Inc.

`lice2cpp` compiler for the `Ice::ServantLocator` interface. Of course, as always, you can add additional member functions, such as a constructor and destructor, and you can add private data members as necessary to support your implementation.

In C++, you can implement the `locate` member function along the following lines:

C++

```
Ice::ObjectPtr
MyServantLocator::locate(const Ice::Current& c, Ice::LocalObjectPtr& cookie)
{
    // Get the object identity. (We use the name member
    // as the database key.)
    //
    std::string name = c.id.name;

    // Use the identity to retrieve the state from the database.
    //
    ServantDetails d;
    try {
        d = DB_lookup(name);
    } catch (const DB_error&)
        return 0;
    }

    // We have the state, instantiate a servant and return it.
    //
    return new PhoneEntryI(d);
}
```

For the time being, the implementations of `finished` and `deactivate` are empty and do nothing.

The `DB_lookup` call in the preceding example is assumed to access the database. If the lookup fails (presumably, because no matching record could be found), `DB_lookup` throws a `DB_error` exception. The code catches that exception and returns zero instead; this raises `ObjectNotExistException` in the client to indicate that the client used a proxy to a no-longer existent Ice object.

Note that `locate` instantiates the servant on the heap and returns it to the Ice run time. This raises the question of when the servant will be destroyed. The answer is that the Ice run time holds onto the servant for as long as necessary, that is, long enough to invoke the operation on the returned servant and to call `finished` once the operation has completed. Thereafter, the servant is no longer needed and the Ice run time destroys the smart pointer that was returned by `locate`. In turn, because no other smart pointers exist for the same servant, this causes the destructor of the `PhoneEntryI` instance to be called, and the servant to be destroyed.

The upshot of this design is that, for every incoming request, we instantiate a servant and allow the Ice run time to destroy the servant once the request is complete. Depending on your application, this may be exactly what is needed, or it may be prohibitively expensive — we will explore designs that avoid creation and destruction of a servant for every request shortly.

In Java, the implementation of our servant locator looks very similar:

Copyright 2016, ZeroC, Inc.

Java

```
public class MyServantLocator implements Ice.ServantLocator {

    public Ice.Object locate(Ice.Current c, Ice.LocalObjectHolder cookie)
    {
        // Get the object identity. (We use the name member
        // as the database key.
        String name = c.id.name;

        // Use the identity to retrieve the state
        // from the database.
        //
        ServantDetails d;
        try {
            d = DB.lookup(name);
        } catch (DB.error e) {
            return null;
        }

        // We have the state, instantiate a servant and return it.
        //
        return new PhoneEntryI(d);
    }

    public void finished(Ice.Current c, Ice.Object servant,
java.lang.Object cookie)
    {
    }

    public void deactivate(String category)
    {
    }
}
```

The C# implementation is virtually identical to the Java implementation, so we do not show it here.

All implementations of `locate` follow the pattern illustrated by the previous pseudo-code:

1. Use the `id` member of the passed `Current` object to obtain the object identity. Typically, only the `name` member of the identity is used to retrieve servant state. The `category` member is normally used to select a servant locator. (We will explore use of the category member shortly.)
2. Retrieve the state of the Ice object from secondary storage (or the network) using the object identity as a key.
   - If the lookup succeeds, you have retrieved the state of the Ice object.
   - If the lookup fails, return null. In that case, the Ice object for the client's request truly does not exist, presumably, because that Ice object was deleted earlier, but the client still has a proxy to the now-deleted object.
3. Instantiate a servant and use the state retrieved from the database to initialize the servant. (In this example, we pass the retrieved state to the servant constructor.)
4. Return the servant.

Of course, before we can use our servant locator, we must inform the adapter of its existence prior to activating the adapter, for example (in Java or C#):

Copyright 2016, ZeroC, Inc.

```
MyServantLocator sl = new MyServantLocator();
adapter.addServantLocator(sl, "");
```

Note that, in this example, we have installed the servant locator for the empty category. This means that `locate` on our servant locator will be called for invocations to any of our Ice objects (because the empty category acts as the default). In effect, with this design, we are not using the `category` member of the object identity. This is fine, as long as all our servants all have the same, single interface. However, if we need to support several different interfaces in the same server, this simple strategy is no longer sufficient.

See Also

- Servant Locators
- Object Identity
- Object Life Cycle
- The Current Object
- Using Identity Categories with Servant Locators

Copyright 2016, ZeroC, Inc.

Using Identity Categories with Servant Locators

Our simple example always instantiates a servant of type `PhoneEntryI`. In other words, the servant locator implicitly is aware of the type of servant the incoming request is for. This is not a very realistic assumption for most servers because, usually, a server provides access to objects with several different interfaces. This poses a problem for our `locate` implementation: somehow, we need to decide inside `locate` what type of servant to instantiate. You have several options for solving this problem:

- Use a separate object adapter for each interface type and use a separate servant locator for each object adapter.

  This technique works fine, but has the down-side that each object adapter requires a separate transport endpoint, which is wasteful.

- Mangle a type identifier into the `name` component of the object identity.

  This technique uses part of the object identity to denote what type of object to instantiate. For example, in our file system application, we have directory and file objects. By convention, we could prepend a 'd' to the identity of every directory and prepend an 'f' to the identity of every file. The servant locator then can use the first letter of the identity to decide what type of servant to instantiate:

  C++

  ```
  Ice::ObjectPtr
  MyServantLocator::locate(const Ice::Current& c, Ice::LocalObjectPtr& cookie)
  {
      // Get the object identity. (We use the name member
      // as the database key.)
      //
      std::string name = c.id.name;
      std::string realId = c.id.name.substr(1);
      try {
          if (name[0] == 'd') {
              // The request is for a directory.
              //
              DirectoryDetails d = DB_lookup(realId);
              return new DirectoryI(d);
          } else {
              // The request is for a file.
              //
              FileDetails d = DB_lookup(realId);
              return new FileI(d);
          }
      } catch (DatabaseNotFoundException&) {
          return 0;
      }
  }
  ```

  While this works, it is awkward: not only do we need to parse the `name` member to work out what type of object to instantiate, but we also need to modify the implementation of `locate` whenever we add a new type to our application.

- Use the `category` member of the object identity to denote the type of servant to instantiate.

  This is the recommended approach: for every interface type, we assign a separate identifier as the value of the `category` member of the object identity. (For example, we can use 'd' for directories and 'f' for files.) Instead of registering a single servant locator, we create two different servant locator implementations, one for directories and one for files, and then register each locator for the appropriate category:

Copyright 2016, ZeroC, Inc.

C++

```
class DirectoryLocator : public virtual Ice::ServantLocator {
public:

    virtual Ice::ObjectPtr locate(const Ice::Current& c,
Ice::LocalObjectPtr& cookie)
    {
        // Code to locate and instantiate a directory here...
    }

    virtual void finished(const Ice::Current& c,
const Ice::ObjectPtr& servant,
                          const Ice::LocalObjectPtr& cookie)
    {
    }

    virtual void deactivate(const std::string& category)
    {
    }
};

class FileLocator : public virtual Ice::ServantLocator {
public:

    virtual Ice::ObjectPtr locate(const Ice::Current& c, Ice::LocalObjectPtr&
cookie)
    {
        // Code to locate and instantiate a file here...
    }

    virtual void finished(const Ice::Current& c,
const Ice::ObjectPtr& servant,
                          const Ice::LocalObjectPtr& cookie)
    {
    }

    virtual void deactivate(const std::string& category)
    {
    }
};

// ...

// Register two locators, one for directories and
// one for files.
//
adapter->addServantLocator(new DirectoryLocator(), "d");
adapter->addServantLocator(new FileLocator(), "f");
```

Yet another option is to use the `category` member of the object identity, but to use a single default servant locator (that is, a locator for the empty category). With this approach, all invocations go to the single default servant locator, and you can switch on the `category` value

Copyright 2016, ZeroC, Inc.

inside the implementation of the `locate` operation to determine which type of servant to instantiate. However, this approach is harder to maintain than the previous one; the `category` member of the Ice object identity exists specifically to support servant locators, so you might as well use it as intended.

See Also

- Servant Locator Example
- Object Identity

Copyright 2016, ZeroC, Inc.

Using Cookies with Servant Locators

Occasionally, it can be useful for a servant locator to pass information between `locate` and `finished`. For example, the implementation of `locate` could choose among a number of alternative database backends, depending on load or availability and, to properly finalize state, the implementation of `finished` might need to know which database was used by `locate`. To support such scenarios, you can create a cookie in your `locate` implementation; the Ice run time passes the value of the cookie to `finished` after the operation invocation has completed. The cookie must derive from `Ice::LocalObject` and can contain whatever state and member functions are useful to your implementation:

C++

```
class MyCookie : public virtual Ice::LocalObject {
public:
    // Whatever is useful here...
};

typedef IceUtil::Handle<MyCookie> MyCookiePtr;

class MyServantLocator : public virtual Ice::ServantLocator {
public:

    virtual Ice::ObjectPtr locate(const Ice::Current& c,
Ice::LocalObjectPtr& cookie)
    {
        // Code as before...

        // Allocate and initialize a cookie.
        //
        cookie = new MyCookie(...);

        return new PhoneEntryI;
    }

    virtual void finished(const Ice::Current& c, const Ice::ObjectPtr& servant,
                          const Ice::LocalObjectPtr& cookie)
    {
        // Down-cast cookie to actual type.
        //
        MyCookiePtr mc = MyCookiePtr::dynamicCast(cookie);

        // Use information in cookie to clean up...
        //
        // ...
    }

    virtual void deactivate(const std::string& category);
};
```

See Also

- Servant Locators

Copyright 2016, ZeroC, Inc.

Default Servants

On this page:

- Overview of Default Servants
- Default Servant API
- Threading Guarantees for Default Servants
- Call Dispatch Semantics for Default Servants
- Guidelines for Implementing Default Servants
  - Object Identity is the Key
  - Minimize Contention
  - Combine Strategies
  - Categories Denote Interfaces
  - Plan for the Future
  - Throw exceptions
  - Handle ice_ping
  - Consider Interceptors
  - Use a Blobject Default Servant to Forward Messages

Overview of Default Servants

The Active Servant Map (ASM) is a simple lookup table that maintains a one-to-one mapping between object identities and servants. Although the ASM is easy to understand and offers efficient indexing, it does not scale well when the number of objects is very large. Scalability is a common problem with object-oriented middleware: servers frequently are used as front ends to large databases that are accessed remotely by clients. The server's job is to present an object-oriented view to clients of a very large number of records in the database. Typically, the number of records is far too large to instantiate servants for even a fraction of the database records.

A common technique for solving this problem is to use *default servants*. A default servant is a servant that, for each request, takes on the persona of a different Ice object. In other words, the servant changes its behavior according to the object identity that is accessed by a request, on a per-request basis. In this way, it is possible to allow clients access to an unlimited number of Ice objects with only a single servant in memory. A default servant is essentially a specialized version of a servant locator that satisfies the majority of use cases with a simpler API, whereas a servant locator provides more flexibility for those applications that require it.

Default servant implementations are attractive not only because of the memory savings they offer, but also because of the simplicity of implementation: in essence, a default servant is a facade [1] to the persistent state of an object in the database. This means that the programming required to implement a default servant is typically minimal: it simply consists of the code required to read and write the corresponding database records.

A default servant is a regular servant that you implement and register with an object adapter. For each incoming request, the object adapter first attempts to locate a servant in its ASM. If no servant is found, the object adapter dispatches the request to a default servant. With this design, a default servant is the object adapter's servant of last resort if no match was found in the ASM.

Implementing a default servant requires a somewhat different mindset than the typical "one servant per Ice object" strategy used in less advanced applications. The most important quality of a default servant is its statelessness: it must be prepared to dispatch multiple requests simultaneously for different objects. The price we have to pay for the unlimited scalability and reduced memory footprint is performance: default servants typically make a database access for every invoked operation, which is obviously slower than caching state in memory as part of a servant that has been added to the ASM. However, this does not mean that default servants carry an unacceptable performance penalty: databases often provide sophisticated caching, so even though the operation implementations read and write the database, as long as they access cached state, performance may be entirely acceptable.

Default Servant API

The default servant API consists of the following operations in the object adapter interface:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local interface ObjectAdapter {
        void addDefaultServant(Object servant, string category);
        Object removeDefaultServant(string category);
        Object findDefaultServant(string category);

        // ...
    };
};
```

As you can see, the object adapter allows you to add, remove, and find default servants. Note that, when you register a default servant, you must provide an argument for the `category` parameter. The value of the `category` parameter controls which object identities the default servant is responsible for: only object identities with a matching `category` member trigger a dispatch to this default servant. An incoming request for which no explicit entry exists in the ASM and with a category for which no default servant is registered returns an `ObjectNotExistException` to the client.

`addDefaultServant` has the following semantics:

- You can register exactly one default servant for a specific category. Attempts to call `addDefaultServant` for the same category more than once raise an `AlreadyRegisteredException`.
- You can register different default servants for different categories, or you can register the same single default servant multiple times (each time for a different category). In the former case, the category is implicit in the default servant instance that is called by the Ice run time; in the latter case, the servant can find out which category the incoming request is for by examining the object identity member of the `Current` object that is passed to the dispatched operation.
- It is legal to register a default servant for the empty category. Such a servant is used if a request comes in for which no entry exists in the ASM, and whose category does not match the category of any other registered default servant.

`removeDefaultServant` removes the default servant for the specified category. Attempts to remove a non-existent default servant raise `NotRegisteredException`. The operation returns the removed default servant. Once a default servant is successfully removed for the specified category, the Ice run time guarantees that no new incoming requests for that category are dispatched to the servant.

The `findDefaultServant` operation allows you to retrieve the default servant for a specific category (including the empty category). If no default servant is registered for the specified category, `findDefaultServant` returns null.

Threading Guarantees for Default Servants

The threading semantics for a default servant are no different than for a servant registered in the ASM: operations may be dispatched on a default servant concurrently, for the same object identity or for different object identities. If you have configured the communicator with multiple dispatch threads, your default servant must protect access to shared data with appropriate locks.

Call Dispatch Semantics for Default Servants

This section summarizes the actions taken by the Ice run time to locate a servant for an incoming request.

Every incoming request implicitly identifies a specific object adapter for the request (because the request arrives at a specific transport endpoint and, therefore, identifies a particular object adapter). The incoming request carries an object identity that must be mapped to a servant. To locate a servant, the Ice run time goes through the following steps, in the order shown:

1. Look for the identity in the ASM. If the ASM contains an entry, dispatch the request to the corresponding servant.
2. If the category of the incoming object identity is non-empty, look for a default servant that is registered for that category. If such a default servant is registered, dispatch the request to that servant.
3. If the category of the incoming object identity is empty, or no default servant could be found for the category in step 2, look for a default servant that is registered for the empty category. If such a default servant is registered, dispatch the request to that servant.
4. If the category of the incoming object identity is non-empty and no servant could be found in the preceding steps, look for a servant locator that is registered for that category. If such a servant locator is registered, call `locate` on the servant locator and, if `locate` returns a servant, dispatch the request to that servant, followed by a call to `finished`; otherwise, if the call to `locate` returns null,

Copyright 2016, ZeroC, Inc.

raise `ObjectNotExistException` or `FacetNotExistException` in the client. (`ObjectNotExistException` is raised if the ASM does not contain a servant with the given identity at all, `FacetNotExistException` is raised if the ASM contains a servant with a matching identity, but a non-matching facet.)
5. If the category of the incoming object identity is empty, or no servant locator could be found for the category in step 4, look for a default servant locator (that is, a servant locator that is registered for the empty category). If a default servant locator is registered, dispatch the request as for step 4.
6. Raise `ObjectNotExistException` or `FacetNotExistException` in the client. (`ObjectNotExistException` is raised if the ASM does not contain a servant with the given identity at all, `FacetNotExistException` is raised if the ASM contains a servant with a matching identity, but a non-matching facet.)

It is important to keep these call dispatch semantics in mind because they enable a number of powerful implementation techniques. Each technique allows you to streamline your server implementation and to precisely control the trade-off between performance, memory consumption, and scalability.

Guidelines for Implementing Default Servants

This section provides some guidelines to assist you in implementing default servants effectively.

Object Identity is the Key

When an incoming request is dispatched to the default servant, the target object identity is provided in the `Current` argument. The `name` fiel d of the identity typically supplies everything the default servant requires in order to satisfy the request. For instance, it may serve as the key in a database query, or even hold an encoded structure in some proprietary format that your application uses to convey more than just a string.

Naturally, the client can also pass arguments to the operation that assist the default servant in retrieving whatever state it requires. However, this approach can easily introduce implementation artifacts into your Slice interfaces, and in most cases the client should not need to know that the server is implemented with a default servant. If at all possible, use only the object identity.

Minimize Contention

For better scalability, the default servant's implementation should strive to eliminate contention among the dispatch threads. As an example, when a database holds the default servant's state, each of the servant's operations usually begins with a query. Assuming that the database API is thread-safe, the servant needs to perform no explicit locking of its own. With a copy of the state in hand, the implementation can work with function-local data to satisfy the request.

Combine Strategies

The ASM still plays a useful role even in applications that are ideally suited for default servants. For example, there is no need to implement a singleton object as a default servant: if there can only be one instance of the object, implementing it as a default servant does nothing to improve your application's scalability.

Applications often install a handful of servants in the ASM while servicing the majority of requests in a default servant. For example, a database application might install a singleton query object in the ASM while using a default servant to process all invocations on the database records.

Categories Denote Interfaces

In general, all of the objects serviced by a default servant must have the same interface. If you only need a default servant for one interface, you can register the default servant with an empty category string. However, to implement several interfaces, you will need a default servant implementation for each one. Furthermore, you must take steps to ensure that the object adapter dispatches an incoming request to the appropriate default servant. The `category` field of the object identity is intended to serve this purpose.

For example, a process control system might have interfaces named `Sensor` and `Switch`. To direct requests to the proper default servant, the application uses the symbol `Sensor` or `Switch` as the category of each object's identity, and registers corresponding default servants having those same categories with the object adapter.

Plan for the Future

If you suspect that you might eventually need to implement more than one interface with default servants, we recommend using a non-empty category even if you start out having only one default servant. Adding another default servant later becomes much easier if the application is already designed to operate correctly with categories.

Throw exceptions

Copyright 2016, ZeroC, Inc.

If a request arrives for an object that no longer exists, it is the default servant's responsibility to raise `ObjectNotExistException` to properly manage object life cycles.

Handle `ice_ping`

One issue you need to be aware of with default servants is the need to override `ice_ping`: the default implementation of `ice_ping` that the servant inherits from its skeleton class always succeeds. For servants that are registered with the ASM, this is exactly what we want; however, for default servants, `ice_ping` must fail if a client uses a proxy to an Ice object that no longer exists. To avoid getting successful `ice_ping` invocations for non-existent Ice objects, you must override `ice_ping` in the default servant. The implementation must check whether the object identity for the request denotes a still-existing Ice object and, if not, raise `ObjectNotExistException`.

It is good practice to override `ice_ping` if you are using default servants. Because you cannot override operations on `Ice::Object` using a Java or C# tie servant (or an Objective-C delegate servant), you must implement default servants by deriving from the generated skeleton class if you choose to override `ice_ping`.

Consider Interceptors

A dispatch interceptor is often installed as a default servant.

Use a `Blobject` Default Servant to Forward Messages

Message forwarding services, such as Glacier2, can be implemented simply and efficiently with a `Blobject` default servant. Such a servant simply chooses a destination to forward a request to, without decoding any of the parameters.

See Also

- The Active Servant Map
- Object Identity
- Servant Locators
- Object Adapters
- The Current Object
- The Ice Threading Model
- Versioning
- Dispatch Interceptors
- Dynamic Invocation and Dispatch Overview
- Glacier2

References

1. Gamma, E., et al. 1994. Design Patterns. Reading, MA: Addison-Wesley.

Copyright 2016, ZeroC, Inc.

Dispatch Interceptors

A dispatch interceptor is a server-side mechanism that allows you to intercept incoming client requests before they are given to a servant. The interceptor can examine the incoming request; in particular, it can see whether the request dispatch is collocation-optimized and examine the `Current` information for the request.

A dispatch interceptor can dispatch a request to a servant and check whether the dispatch was successful; if not, the interceptor can choose to retry the dispatch. This functionality is useful to automatically retry requests that have failed due to a recoverable error condition, such as a database deadlock exception. (Freeze uses dispatch interceptors for this purpose in its evictor implementations.)

On this page:

- Dispatch Interceptor API
  - Objective-C Mapping for Dispatch Interceptors
- Using a Dispatch Interceptor

Dispatch Interceptor API

Dispatch interceptors are not defined in Slice, but are provided as an API that is specific to each programming language. The remainder of this section presents the interceptor API for C++; for Java and .NET, the API is analogous, so we do not show it here.

In C++, a dispatch interceptor has the following interface:

C++

```
namespace Ice {
    class DispatchInterceptor : public virtual Object {
    public:
        virtual DispatchStatus dispatch(Request&) = 0;
    };

    typedef IceInternal::Handle<DispatchInterceptor> DispatchInterceptorPtr;
}
```

Note that a `DispatchInterceptor` *is-a* `Object`, that is, you use a dispatch interceptor as a servant.

To create a dispatch interceptor, you must derive a class from `DispatchInterceptor` and provide an implementation of the pure virtual `dispatch` function. The job of `dispatch` is to pass the request to the servant and to return a dispatch status, defined as follows:

C++

```
namespace Ice {
    enum DispatchStatus {
        DispatchOK, DispatchUserException, DispatchAsync
    };
}
```

The enumerators indicate how the request was dispatched:

- `DispatchOK`
  The request was dispatched synchronously and completed without an exception.
- `DispatchUserException`
  The request was dispatched synchronously and raised a user exception.
- `DispatchAsync`

Copyright 2016, ZeroC, Inc.

The request was dispatched successfully as an asynchronous request; the result of the request is not available to the interceptor because the result is delivered to the AMD callback when the request completes.

The Ice run time provides basic information about the request to the `dispatch` function in the form of a `Request` object:

C++

```
namespace Ice {
    class Request {
    public:
        virtual const Current& getCurrent();
    };
}
```

- `getCurrent` provides access to the `Current` object for the request, which provides access to information about the request, such as the object identity of the target object, the object adapter used to dispatch the request, and the operation name.

Note that `Request`, for performance reasons, is *not* thread-safe. This means that you must not concurrently dispatch from different threads using the same `Request` object. (Concurrent dispatch for different requests does not cause any problems.)

To use a dispatch interceptor, you instantiate your derived class and register it as a servant with the Ice run time in the usual way, such as by adding the interceptor to the Active Servant Map (ASM), or returning the interceptor as a servant from a call to `locate` on a servant locator.

Objective-C Mapping for Dispatch Interceptors

The Objective-C mapping in Ice Touch does not support AMD, therefore the return type of the `dispatch` method is simplified to a boolean:

Objective-C

```
@protocol ICEDispatchInterceptor <ICEObject>
-(BOOL) dispatch:(id<ICERequest>)request;
@end
```

A return value of `YES` is equivalent to `DispatchOK` and indicates that the request completed without an exception. A return value of `NO` is equivalent to `DispatchUserException`.

Using a Dispatch Interceptor

Your implementation of the `dispatch` function must dispatch the request to the actual servant. Here is a very simple example implementation of an interceptor that dispatches the request to the servant passed to the interceptor's constructor:

Copyright 2016, ZeroC, Inc.

C++

```
class InterceptorI : public Ice::DispatchInterceptor {
public:
    InterceptorI(const Ice::ObjectPtr& servant)
        : _servant(servant) {}

    virtual Ice::DispatchStatus dispatch(Ice::Request& request) {
        return _servant->ice_dispatch(request);
    }

    Ice::ObjectPtr _servant;
};
```

Note that our implementation of `dispatch` calls `ice_dispatch` on the target servant to dispatch the request. `ice_dispatch` does the work of actually (synchronously) invoking the operation.

Also note that `dispatch` returns whatever is returned by `ice_dispatch`. For synchronous dispatch, you should always implement your interceptor in this way and not change this return value.

We can use this interceptor to intercept requests to a servant of any type as follows:

C++

```
ExampleIPtr servant = new ExampleI;
Ice::DispatchInterceptorPtr interceptor = new InterceptorI(servant);
adapter->add(interceptor, communicator->stringToIdentity("ExampleServant"));
```

Note that, because dispatch interceptor *is-a* servant, this means that the servant to which the interceptor dispatches need not be the actual servant. Instead, it could be another dispatch interceptor that ends up dispatching to the real servant. In other words, you can chain dispatch interceptors; each interceptor's `dispatch` function is called until, eventually, the last interceptor in the chain dispatches to the actual servant.

A more interesting use of a dispatch interceptor is to retry a call if it fails due to a recoverable error condition. Here is an example that retries a request if it raises a local exception defined in Slice as follows:

Slice

```
local exception DeadlockException { /* ... */ };
```

Note that this is a `local` exception. Local exceptions that are thrown by the servant propagate to `dispatch` and can be caught there. A database might throw such an exception if the database detects a locking conflict during an update. We can retry the request in response to this exception using the following `dispatch` implementation:

Copyright 2016, ZeroC, Inc.

C++

```
virtual Ice::DispatchStatus dispatch(Ice::Request& request) {
    while (true) {
        try {
            return _servant->ice_dispatch(request);
        } catch (const DeadlockException&) {
            // Happens occasionally
        }
    }
}
```

Of course, a more robust implementation might limit the number of retries and possibly add a delay before retrying.

You can also retry an asynchronous dispatch. In this case, each asynchronous dispatch attempt creates a new AMD callback object.

- If the response for the retried request has been sent already, the interceptor receives a `ResponseSentException`. Your interceptor must either not handle this exception (or rethrow it) or return `DispatchAsync`.
- If the response for the request has not been sent yet, the Ice run time ignores any call to `ice_response` or `ice_exception` on the old AMD callback.

If an operation throws a user exception (as opposed to a local exception), the user exception cannot be caught by `dispatch` as an exception but, instead, is reported by the return value of `ice_dispatch`: a return value of `DispatchUserException` indicates that the operation raised a user exception. You can retry a request in response to a user exception as follows:

C++

```
virtual Ice::DispatchStatus dispatch(Ice::Request& request) {
    Ice::DispatchStatus d;
    do {
        d = _servant->ice_dispatch(request);
    } while (d == Ice::DispatchUserException);
    return d;
}
```

This is fine as far as it goes, but not particularly useful because the preceding code retries if *any* kind of user exception is thrown. However, typically, we want to retry a request only if a *specific* user exception is thrown. The problem here is that the `dispatch` function does not have direct access to the actual exception that was thrown — all it knows is that *some* user exception was thrown, but not which one.

To retry a request for a specific user exception, you need to implement your servants such that they leave some "footprint" behind if they throw the exception of interest. This allows your request interceptor to test whether the user exception should trigger a retry. There are various techniques you can use to achieve this. For example, you can use thread-specific storage to test a retry flag that is set by the servant if it throws the exception or, if you use transactions, you can attach the retry flag to the transaction context. However, doing so is more complex; the intended use case is to permit retry of requests in response to local exceptions, so we suggest you retry requests only for local exceptions.

The most common use case for a dispatch interceptor is as a default servant. Rather than having an explicit interceptor for individual servants, you can register a dispatch interceptor as default servant. You can then choose the "real" servant to which to dispatch the request inside `dispatch`, prior to calling `ice_dispatch`. This allows you to intercept and selectively retry requests based on their outcome, which cannot be done using a servant locator.

See Also

- The Current Object
- Collocated Invocation and Dispatch

Copyright 2016, ZeroC, Inc.

- The Active Servant Map
- Servant Locators
- Freeze Evictors
- Default Servants

Copyright 2016, ZeroC, Inc.

Client-Server Features

This section presents APIs and features best understood while considering both sides of a client-server interaction.

Client-specific and server-specific features are presented in separate sections, Client-Side Features and Server-Side Features.

Topics

- The Ice Threading Model
- Connection Management
- Connection Timeouts
- Collocated Invocation and Dispatch
- Locators
- Slicing Values and Exceptions
- Dynamic Ice
- Facets
- Versioning

Copyright 2016, ZeroC, Inc.

The Ice Threading Model

Ice is inherently a multi-threaded platform. There is no such thing as a single-threaded server in Ice. As a result, you must concern yourself with concurrency issues: if a thread reads a data structure while another thread updates the same data structure, havoc will ensue unless you protect the data structure with appropriate locks. In order to build Ice applications that behave correctly, it is important that you understand the threading semantics of the Ice run time. Here we discuss Ice's *thread pool* concurrency model and provide guidelines for writing thread-safe Ice applications.

Topics

- Thread Pools
- Object Adapter Thread Pools
- Thread Pool Design Considerations
- Concurrent Proxy Invocations
- Nested Invocations
- Thread Safety
- Dispatching Invocations to User Threads
- Blocking API Calls

Copyright 2016, ZeroC, Inc.

Thread Pools

A thread pool is a collection of threads that the Ice run time draws upon to perform specific tasks.

On this page:

- Introduction to Thread Pools
- Configuring Thread Pools
- Dynamic Thread Pools

Introduction to Thread Pools

Each communicator creates two thread pools:

- The *client thread pool* services outgoing connections, which primarily involves handling the replies to outgoing requests and includes notifying AMI callback objects. If a connection is used in bidirectional mode, the client thread pool also dispatches incoming callback requests.
- The *server thread pool* services incoming connections. It dispatches incoming requests and, for bidirectional connections, processes replies to outgoing requests.

By default, these two thread pools are shared by all of the communicator's object adapters. If necessary, you can configure individual object adapters to use a private thread pool instead.

If a thread pool is exhausted because all threads are currently dispatching a request, additional incoming requests are transparently delayed until a request completes and relinquishes its thread; that thread is then used to dispatch the next pending request. Ice minimizes thread context switches in a thread pool by using a leader-follower implementation [1].

Configuring Thread Pools

Each thread pool has a unique name that serves as the prefix for its configuration properties:

- *name*`.Size`
  This property specifies the initial size of the thread pool. If not defined, the default value is 1.
- *name*`.SizeMax`
  This property specifies the maximum size of the thread pool. If not defined, the default value is 1. If the value of this property is less than that of *name*`.Size`, this property is adjusted to be equal to *name*`.Size`.
- *name*`.SizeWarn`
  This property sets a high water mark; when the number of threads in a pool reaches this value, the Ice run time logs a warning message. If you see this warning message frequently, it could indicate that you need to increase the value of *name*`.SizeMax`. The default value is 0, which disables the warning.
- *name*`.StackSize`
  This property specifies the number of bytes to use as the stack size of threads in the thread pool. The operating system's default is used if this property is not defined or is set to 0.
- *name*`.Serialize`
  Setting this property to a value greater than 0 forces the thread pool to serialize all messages received over a connection. It is unnecessary to enable serialization for a thread pool whose maximum size is 1 because such a thread pool is already limited to processing one message at a time. For thread pools with more than one thread, serialization has a negative impact on latency and throughput. If not defined, the default value is 0. We discuss this feature in more detail in Thread Pool Design Considerations.
- *name*`.ThreadIdleTime`
  This property specifies the number of seconds that a thread in the thread pool must be idle before it terminates. The default value is 60 seconds if this property is not defined. Setting it to 0 disables the termination of idle threads.

For configuration purposes, the names of the client and server thread pools are `Ice.ThreadPool.Client` and `Ice.ThreadPool.Server`, respectively. As an example, the following properties establish the initial and maximum sizes for these thread pools:

```
Ice.ThreadPool.Client.Size=1
Ice.ThreadPool.Client.SizeMax=10
Ice.ThreadPool.Server.Size=1
Ice.ThreadPool.Server.SizeMax=10
```


Copyright 2016, ZeroC, Inc.

To monitor the thread pool activities of the Ice run time, you can enable the `Ice.Trace.ThreadPool` property. Setting this property to a non-0 value causes the Ice run time to log a message when it creates a thread pool, as well as each time the size of a thread pool increases or decreases.

Dynamic Thread Pools

A *dynamic* thread pool can grow and shrink when necessary in response to changes in an application's work load. All thread pools have at least one thread, but a dynamic thread pool can grow as the demand for threads increases, up to the pool's maximum size. Threads may also be terminated automatically when they have been idle for some time.

The dynamic nature of a thread pool is determined by the configuration properties *name*`.Size`, *name*`.SizeMax`, and *name*`.ThreadIdleTime`. A thread pool is not dynamic in its default configuration because *name*`.Size` and *name*`.SizeMax` are both set to 1, meaning the pool can never grow to contain more than a single thread. To configure a dynamic thread pool, you must set at least one of *name*`.Size` or *name*`.SizeMax` to a value greater than 1. We can use several configuration scenarios to explore the semantics of dynamic thread pools in greater detail:

- *name*`.SizeMax=5`

  This thread pool initially contains a single thread because *name*`.Size` has a default value of 1, and Ice can grow the pool up to the maximum of 5 threads. During periods of inactivity, idle threads terminate after 60 seconds (the default value for *name*`.ThreadIdleTime`) until the pool contains just 1 thread again.

- *name*`.Size=3`
  name`.SizeMax=5`

  This thread pool starts with 3 active threads but otherwise behaves the same as in the previous configuration. The pool can still shrink to a size of 1 as threads become idle.

- *name*`.Size=3`
  name`.ThreadIdleTime=10`

  This thread pool starts with 3 active threads and shrinks quickly to 1 thread during periods of inactivity. As demand increases again, the thread pool can return to its maximum size of 3 threads (*name*`.SizeMax` defaults to the value of *name*`.Size`).

- *name*`.SizeMax=5`
  name`.ThreadIdleTime=0`

  This thread pool can grow from its initial size of 1 thread to contain up to 5 threads, but it will never shrink because *name*`.ThreadIdleTime` is set to 0.

- *name*`.Size=5`
  name`.ThreadIdleTime=0`

  This thread pool starts with 5 threads and can neither grow nor shrink.

To summarize, the value of *name*`.ThreadIdleTime` determines whether (and how quickly) a thread pool can shrink to a size of 1. A thread pool that shrinks can also grow to its maximum size. Finally, setting *name*`.SizeMax` to a value larger than *name*`.Size` allows a thread pool to grow beyond its initial capacity.

See Also

- Thread Pool Design Considerations
- Bidirectional Connections
- Object Adapters
- Object Adapter Thread Pools
- Ice.ThreadPool.*

References

1. Schmidt, D. C. et al. 2000. "Leader/Followers: A Design Pattern for Efficient Multi-Threaded Event Demultiplexing and Dispatching". In *Proceedings of the 7th Pattern Languages of Programs Conference*, WUCS-00-29, Seattle, WA: University of Washington.

Copyright 2016, ZeroC, Inc.

Object Adapter Thread Pools

The default behavior of an object adapter is to share the thread pools of its communicator and, for many applications, this behavior is entirely sufficient. However, the ability to configure an object adapter with its own thread pool is useful in certain situations:

- When the concurrency requirements of an object adapter does not match those of its communicator. In a server with multiple object adapters, the configuration of the communicator's client and server thread pools may be a good match for some object adapters, but others may have different requirements. For example, the servants hosted by one object adapter may not support concurrent access, in which case limiting that object adapter to a single-threaded pool eliminates the need for synchronization in those servants. On the other hand, another object adapter might need a multi-threaded pool for better performance.
- To ensure that a minimum number of threads is available for dispatching requests to an adapter's servants. This is especially important for eliminating the possibility of deadlocks when using nested invocations.

An object adapter's thread pool supports all of the properties described in Configuring Thread Pools. For configuration purposes, the name of an adapter's thread pool is *`adapter`*`.ThreadPool`, where *`adapter`* is the name of the adapter.

An adapter creates its own thread pool when at least one of the following properties has a value greater than zero:

- *`adapter`*`.ThreadPool.Size`
- *`adapter`*`.ThreadPool.SizeMax`

These properties have the same semantics as those described earlier except they both have a default value of zero, meaning that an adapter uses the communicator's thread pools by default.

As an example, the properties shown below configure a thread pool for the object adapter named `PrinterAdapter`:

```
PrinterAdapter.ThreadPool.Size=3
PrinterAdapter.ThreadPool.SizeMax=15
PrinterAdapter.ThreadPool.SizeWarn=14
```

See Also

- Thread Pools
- Object Adapters
- Nested Invocations
- Object Adapter Properties

Copyright 2016, ZeroC, Inc.

Thread Pool Design Considerations

Improper configuration of a thread pool can have a serious impact on the performance of your application. This page discusses some issues that you should consider when designing and configuring your applications.

On this page:

- Single-Threaded Pool
- Multi-Threaded Pool
- Serializing Requests in a Multi-Threaded Pool

Single-Threaded Pool

There are several implications of using a thread pool with a maximum size of one thread:

- **Only one message can be dispatched at a time.**

  This can be convenient because it lets you avoid (or postpone) dealing with thread-safety issues in your application. However, it also eliminates the possibility of dispatching requests concurrently, which can be a bottleneck for applications running on multi-CPU systems or that perform blocking operations. Another option is to enable serialization in a multi-threaded pool.

- **Only one AMI reply can be processed at a time.**

  An application must increase the size of the client thread pool in order to process multiple AMI callbacks in parallel.

- **Nested twoway invocations are limited.**

  At most one level of nested twoway invocations is possible.

It is important to remember that a communicator's client and server thread pools have a default maximum size of one thread, therefore these limitations also apply to any object adapter that shares the communicator's thread pools.

Multi-Threaded Pool

Configuring a thread pool to support multiple threads implies that the application is prepared for the Ice run time to dispatch operation invocations or AMI callbacks concurrently. Although greater effort is required to design a thread-safe application, you are rewarded with the ability to improve the application's scalability and throughput.

Choosing an appropriate maximum size for a thread pool requires careful analysis of your application. For example, in compute-bound applications it is best to limit the number of threads to the number of physical processor cores or threads on the host machine; adding any more threads only increases context switches and reduces performance. Increasing the size of the pool beyond the number of cores can improve responsiveness when threads can become blocked while waiting for the operating system to complete a task, such as a network or file operation. On the other hand, a thread pool configured with too many threads can have the opposite effect and negatively impact performance. Testing your application in a realistic environment is the recommended way of determining the optimum size for a thread pool.

If your application uses nested invocations, it is very important that you evaluate whether it is possible for thread starvation to cause a deadlock. Increasing the size of a thread pool can lessen the chance of a deadlock, but other design solutions are usually preferred.

Serializing Requests in a Multi-Threaded Pool

When using a multi-threaded pool, the nondeterministic nature of thread scheduling means that requests from the same connection may not be dispatched in the order they were received. Some applications cannot tolerate this behavior, such as a transaction processing server that must guarantee that requests are executed in order. There are two ways of satisfying this requirement:

1. Use a single-threaded pool.
2. Configure a multi-threaded pool to serialize requests using its `Serialize` property.

At first glance these two options may seem equivalent, but there is a significant difference: a single-threaded pool can only dispatch one request at a time and therefore serializes requests from *all* connections, whereas a multi-threaded pool configured for serialization can dispatch requests from different connections concurrently while serializing requests from the same connection.

You can obtain the same behavior from a multi-threaded pool without enabling serialization, but only if you design the clients so that they do not send requests from multiple threads, do not send requests over more than one connection, and only use synchronous twoway invocations. In general, however, it is better to avoid such tight coupling between the implementations of the client and server.

Enabling serialization can improve responsiveness and performance compared to a single-threaded pool, but there is an associated cost. The extra synchronization that the pool must perform to serialize requests adds significant overhead and results in higher latency and reduced throughput.

Copyright 2016, ZeroC, Inc.

As you can see, thread pool serialization is not a feature that you should enable without analyzing whether the benefits are worthwhile. For example, it might be an inappropriate choice for a server with long-running operations when the client needs the ability to have several operations in progress simultaneously. If serialization was enabled in this situation, the client would be forced to work around it by opening several connections to the server, which again tightly couples the client and server implementations. If the server must keep track of the order of client requests, a better solution would be to use serialization in conjunction with asynchronous dispatch to queue the incoming requests for execution by other threads.

See Also

- Thread Pools
- Concurrent Proxy Invocations
- Nested Invocations
- Thread Safety
- Connection Establishment

Copyright 2016, ZeroC, Inc.

Concurrent Proxy Invocations

Proxy objects are fully thread safe, meaning a client can invoke on the same proxy object from multiple threads concurrently without the need for additional synchronization. However, Ice makes few guarantees about the order in which it sends proxy invocations over a connection.

To understand the ordering issue, it's important to first understand some fundamental proxy concepts:

- At any point in time, a proxy may or may not be associated with a connection.
- A new proxy initially has no connection, and Ice does not attempt to associate it with a connection until its first invocation.
- If Ice needs to establish a new connection for a proxy, Ice queues all invocations on that proxy until the connection succeeds.
- After a proxy is associated with a connection, the proxy may or may not cache that connection for subsequent invocations.
- Proxies share connections by default, but an application can force proxies to use separate connections.

Ice (as of version 3.6) guarantees that ordering will be maintained for invocations on the same proxy object, but only if that proxy caches its connection. This guarantee also holds for invocations on two or more proxies that programmatically compare as equal, meaning the proxies have the same identity and other configuration settings, and furthermore these proxies must cache connections as well.

Ice does not guarantee the order of invocations for any other situation.

The order in which the Ice run time in a client sends invocations over a connection does not necessarily determine the order in which they will be executed in the server. See Thread Pool Design Considerations for information about ordering guarantees in servers.

See Also

- Proxies
- Connection Establishment
- Thread Pool Design Considerations

Copyright 2016, ZeroC, Inc.

Nested Invocations

A *nested invocation* is one that is made within the context of another Ice operation. For instance, the implementation of an operation in a servant might need to make a nested invocation on some other object, or an AMI callback object might invoke an operation in the course of processing a reply to an asynchronous request. It is also possible for one of these invocations to result in a nested callback to the originating process. The maximum depth of such invocations is determined by the size of the thread pools used by the communicating parties.

On this page:

- Deadlocks with Nested Invocations
- Analyzing an Application for Nested Invocations

Deadlocks with Nested Invocations

Applications that use nested invocations must be carefully designed to avoid the potential for deadlock, which can easily occur when invocations take a circular path. For example, this illustration presents a deadlock scenario when using the default thread pool configuration:

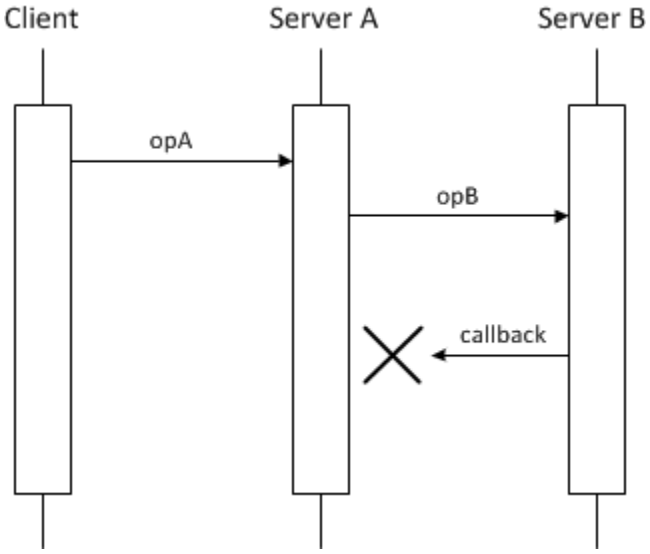

Nested invocation deadlock.

In this diagram, the implementation of `opA` makes a nested twoway invocation of `opB`, but the implementation of `opB` causes a deadlock when it tries to make a nested callback. As mentioned in Thread Pools, the communicator's thread pools have a maximum size of one thread unless explicitly configured otherwise. In Server A, the only thread in the server thread pool is busy waiting for its invocation of `opB` to complete, and therefore no threads remain to handle the callback from Server B. The client is now blocked because Server A is blocked, and they remain blocked indefinitely unless timeouts are used.

There are several ways to avoid a deadlock in this scenario:

- **Increase the maximum size of the server thread pool in Server A.**

  Configuring the server thread pool in Server A to support more than one thread allows the nested callback to proceed. This is the simplest solution, but it requires that you know in advance how deeply nested the invocations may occur, or that you set the maximum size to a sufficiently large value that exhausting the pool becomes unlikely. For example, setting the maximum size to two avoids a deadlock when a single client is involved, but a deadlock could easily occur again if multiple clients invoke `opA` simultaneou sly. Furthermore, setting the maximum size too large can cause its own set of problems.

- **Use a oneway invocation.**

  If Server A called `opB` using a oneway invocation, it would no longer need to wait for a response and therefore `opA` could complete, making a thread available to handle the callback from Server B. However, we have made a significant change in the semantics of o pA because now there is no guarantee that `opB` has completed before `opA` returns, and it is still possible for the oneway invocation of `opB` to block.

- **Create another object adapter for the callbacks.**

  No deadlock occurs if the callback from Server B is directed to a different object adapter that is configured with its own thread pool.

- **Implement `opA` using asynchronous dispatch and invocation.**

Copyright 2016, ZeroC, Inc.

By declaring `opA` as an AMD operation and invoking `opB` using AMI, Server A can avoid blocking the thread pool's thread while it waits for `opB` to complete. This technique, known as *asynchronous request chaining*, is used extensively in Ice services such as IceGrid and Glacier2 to eliminate the possibility of deadlocks.

As another example, consider a client that makes a nested invocation from an AMI callback object using the default thread pool configuration. The (one and only) thread in the client thread pool receives the reply to the asynchronous request and invokes its callback object. If the callback object in turn makes a nested twoway invocation, a deadlock occurs because no more threads are available in the client thread pool to process the reply to the nested invocation. The solutions are similar to some of those presented in the above illustration: increase the maximum size of the client thread pool, use a oneway invocation, or call the nested invocation using AMI.

Analyzing an Application for Nested Invocations

A number of factors must be considered when evaluating whether an application is properly designed and configured for nested invocations:

- The thread pool configurations in use by all communicating parties have a significant impact on an application's ability to use nested invocations. While analyzing the path of circular invocations, you must pay careful attention to the threads involved to determine whether sufficient threads are available to avoid deadlock. This includes not just the threads that dispatch requests, but also the threads that make the requests and process the replies. Enabling the `Ice.Trace.ThreadPool` property can give you a better understanding of the thread pool behavior in your application.
- Bidirectional connections are another complication, since you must be aware of which threads are used on either end of the connection.
- Finally, the synchronization activities of the communicating parties must also be scrutinized. For example, a deadlock is much more likely when a lock is held while making an invocation.

As you can imagine, tracing the call flow of a distributed application to ensure there is no possibility of deadlock can quickly become a complex and tedious process. In general, it is best to avoid circular invocations if at all possible.

See Also

- Thread Pools
- Object Adapter Thread Pools
- Thread Pool Design Considerations
- Oneway Invocations

Copyright 2016, ZeroC, Inc.

Thread Safety

The Ice run time itself is fully thread safe, meaning multiple application threads can safely call methods on objects such as communicators, object adapters, and proxies without synchronization problems. As a developer, you must also be concerned with thread safety because the Ice run time can dispatch multiple invocations concurrently in a server. In fact, it is possible for multiple requests to proceed in parallel within the same servant and within the same operation on that servant. It follows that, if the operation implementation manipulates non-stack storage (such as member variables of the servant or global or static data), you must interlock access to this data to avoid data corruption.

The need for thread safety in an application depends on its configuration. Using the default thread pool configuration typically makes synchronization unnecessary because at most one operation can be dispatched at a time. Thread safety becomes an issue once you increase the maximum size of a thread pool.

Ice uses the native synchronization and threading primitives of each platform. For C++ users, Ice provides a collection of convenient and portable wrapper classes for use by Ice applications.

On this page:

- Threading Issues with Marshaling
- Thread Creation and Destruction Hooks
- Installing Thread Hooks with a Plug-in

Threading Issues with Marshaling

The marshaling semantics of the Ice run time present a subtle thread safety issue that arises when an operation returns data by reference. In C++, the only relevant case is returning an instance of a Slice class, either directly or nested as a member of another type. In Java, .NET, and the scripting languages, Slice structures, sequences, and dictionaries are also affected.

The potential for corruption occurs whenever a servant returns data by reference, yet continues to hold a reference to that data. For example, consider the following Java implementation:

Java

```
public class GridI extends _GridDisp
{
    GridI()
    {
        _grid = // ...
    }

    public int[][]
    getGrid(Ice.Current curr)
    {
        return _grid;
    }

    public void
    setValue(int x, int y, int val, Ice.Current curr)
    {
        _grid[x][y] = val;
    }

    private int[][] _grid;
}
```

Suppose that a client invoked the `getGrid` operation. While the Ice run time marshals the returned array in preparation to send a reply message, it is possible for another thread to dispatch the `setValue` operation on the same servant. This race condition can result in several unexpected outcomes, including a failure during marshaling or inconsistent data in the reply to `getGrid`. Synchronizing the `getGrid` and `setValue` operations would not fix the race condition because the Ice run time performs its marshaling outside of this synchronization.

Copyright 2016, ZeroC, Inc.

One solution is to implement accessor operations, such as `getGrid`, so that they return copies of any data that might change. There are several drawbacks to this approach:

- Excessive copying can have an adverse affect on performance.
- The operations must return deep copies in order to avoid similar problems with nested values.
- The code to create deep copies is tedious and error-prone to write.

Another solution is to make copies of the affected data only when it is modified. In the revised code shown below, `setValue` replaces `_grid` with a copy that contains the new element, leaving the previous contents of `_grid` unchanged:

Java

```
public class GridI extends _GridDisp
{
    ...

    public synchronized int[][]
    getGrid(Ice.Current curr)
    {
        return _grid;
    }

    public synchronized void
    setValue(int x, int y, int val, Ice.Current curr)
    {
        int[][] newGrid = // shallow copy...
        newGrid[x][y] = val;
        _grid = newGrid;
    }

    ...
}
```

This allows the Ice run time to safely marshal the return value of `getGrid` because the array is never modified again. For applications where data is read more often than it is written, this solution is more efficient than the previous one because accessor operations do not need to make copies. Furthermore, intelligent use of shallow copying can minimize the overhead in mutating operations.

Finally, a third approach changes accessor operations to use AMD in order to regain control over marshaling. After annotating the `getGrid` operation with `amd` metadata, we can revise the servant as follows:

Copyright 2016, ZeroC, Inc.

Java

```
public class GridI extends _GridDisp
{
    ...

    public synchronized void
    getGrid_async(AMD_Grid_getGrid cb, Ice.Current curr)
    {
        cb.ice_response(_grid);
    }

    public synchronized void
    setValue(int x, int y, int val, Ice.Current curr)
    {
        _grid[x][y] = val;
    }

    ...
}
```

Normally, AMD is used in situations where the servant needs to delay its response to the client without blocking the calling thread. For `getGrid`, that is not the goal; instead, as a side-effect, AMD provides the desired marshaling behavior. Specifically, the Ice run time marshals the reply to an asynchronous request at the time the servant invokes `ice_response` on the AMD callback object. Because `getGrid` and `setValue` are synchronized, this guarantees that the data remains in a consistent state during marshaling.

Thread Creation and Destruction Hooks

On occasion, it is necessary to intercept the creation and destruction of threads created by the Ice run time, for example, to interoperate with libraries that require applications to make thread-specific initialization and finalization calls (such as COM's `CoInitializeEx` and `CoUninitialize`). Ice provides callbacks to inform an application when each run-time thread is created and destroyed. For C++, the callback class looks as follows:

C++

```
class ThreadNotification : public IceUtil::Shared {
public:
    virtual void start() = 0;
    virtual void stop() = 0;
};
typedef IceUtil::Handle<ThreadNotification> ThreadNotificationPtr;
```

To receive notification of thread creation and destruction, you must derive a class from `ThreadNotification` and implement the `start` and `stop` member functions. These functions will be called by the Ice run by each thread as soon as it is created, and just before it exits. You must install your callback class in the Ice run time when you create a communicator by setting the `threadHook` member of the `InitializationData` structure.

For example, you could define a callback class and register it with the Ice run time as follows:

Copyright 2016, ZeroC, Inc.

C++

```
class MyHook : public virtual Ice::ThreadNotification {
public:
    void start()
    {
        cout << "start: id = " << ThreadControl().id() << endl;
    }
    void stop()
    {
        cout << "stop: id = " << ThreadControl().id() << endl;
    }
};

int
main(int argc, char* argv[])
{
    // ...

    Ice::InitializationData id;
    id.threadHook = new MyHook;
    communicator = Ice::initialize(argc, argv, id);

    // ...
}
```

The implementation of your `start` and `stop` methods can make whatever thread-specific calls are required by your application.

For Java and C#, `Ice.ThreadNotification` is an interface:

```
public interface ThreadNotification {
    void start();
    void stop();
}
```

To receive the thread creation and destruction callbacks, you must derive a class from this interface that implements the `start` and `stop` methods, and register an instance of that class when you create the communicator. (The code to do this is analogous to the C++ version.)

For Python, the interface is:

Copyright 2016, ZeroC, Inc.

Python

```
class ThreadNotification(object):
    def __init__(self):
        pass

    # def start():
    # def stop():
```

The Ice run time calls the `start` and `stop` methods of the class instance you provide to `Ice.initialize` when it creates and destroys threads.

Installing Thread Hooks with a Plug-in

The thread hook facility described above requires that you modify a program's source code in order to receive callbacks when threads in the Ice run time are created and destroyed. It is also possible to install thread hooks using the Ice plug-in facility, which is useful for adding thread hooks to an existing program that you cannot (or prefer not to) modify.

Ice provides a base class named `ThreadHookPlugin` for C++, Java, and C# that supplies the necessary functionality. The C++ class definition is shown below:

C++

```
namespace Ice {
class ThreadHookPlugin : public Ice::Plugin {
public:

    ThreadHookPlugin(const CommunicatorPtr& communicator,
const ThreadNotificationPtr&);

    virtual void initialize();

    virtual void destroy();
};
}
```

The equivalent definitions for Java and C# are quite similar and therefore not presented here.

The `ThreadHookPlugin` constructor installs the given `ThreadNotification` object into the specified communicator. The `initialize` a nd `destroy` methods are empty, but you can subclass `ThreadHookPlugin` and override these methods if necessary.

In order to create a thread hook plug-in, you must do the following:

- Define and export a factory class (for Java and C#) or factory function (for C++) that returns an instance of `ThreadHookPlugin`, as described in the plug-in API.
- Implement the `ThreadNotification` object that you will pass to the `ThreadHookPlugin` constructor.
- Package your code into a format that is suitable for dynamic loading, such as a shared library or DLL for C++ or an assembly for C#.

To install your plug-in, use a configuration property like the one shown below:

Copyright 2016, ZeroC, Inc.

```
Ice.Plugin.MyThreadHookPlugin=MyHooks:createPlugin ...
```

The first component of the property value represents the plug-in's entry point. For C++, this value includes the abbreviated name of the shared library or DLL (`MyHooks`) and the name of a factory function (`createPlugin`).

If your property value is language-specific and the configuration file containing this property is shared by programs in multiple implementation languages, you can use an alternate syntax that is loaded only by the Ice run time for a certain language. For example, here is the C++-specific version:

```
Ice.Plugin.MyThreadHookPlugin.cpp=MyHooks:createPlugin ...
```

For more information, see Ice.Plugin.*.

See Also

- Communicator Initialization
- Threads and Concurrency with C++
- Plug-in Facility
- Plug-in API
- Ice.Plugin.*
- Ice.InitPlugins
- Ice.PluginLoadOrder

Copyright 2016, ZeroC, Inc.

Dispatching Invocations to User Threads

By default, operation invocations and AMI callbacks are executed by a thread from a thread pool. This behavior is simple and convenient for applications because they need not concern themselves with thread creation and destruction. However, there are situations where it is necessary to respond to operation invocations or AMI callbacks in a particular thread. For example, in a server, you might need to update a database that does not permit concurrent access from different threads or, in a client, you might need to update a user interface with the results of an invocation. (Many UI frameworks require all UI updates to be made by a specific thread.)

In Ice for C++, Java, .NET, and Objective-C, you can control which thread receives operation invocations and AMI callbacks, so you can ensure that all updates are made by a thread you choose. The implementation techniques vary slightly for each language and are explained in the sections that follow.

On this page:

- C++ Dispatcher API
  - Dispatching with C++11 Lambda Functions
- Java Dispatcher API
- C# Dispatcher API
- Objective-C Dispatcher API
- Dispatcher Implementation Notes

C++ Dispatcher API

To install a dispatcher, you must instantiate a class that derives from `Ice::Dispatcher` and initialize a communicator with that instance in the `InitializationData` structure. All invocations that arrive for this communicator are made via the specified dispatcher. For example:

C++

```
class MyDispatcher : public Ice::Dispatcher /*, ... */
    // ...
};

int
main(int argc, char* argv[])
{
    Ice::CommunicatorPtr communicator;

    try {
        Ice::InitializationData initData;
        initData.properties = Ice::createProperties(argc, argv);
        initData.dispatcher = new MyDispatcher();
        communicator = Ice::initialize(argc, argv, initData);

        // ...
    } catch (const Ice::Exception& ex) {
        // ...
    }

    // ...
}
```

The `Ice::Dispatcher` abstract base class has the following interface:

Copyright 2016, ZeroC, Inc.

C++

```
class Dispatcher : virtual public IceUtil::Shared
{
public:
    virtual void dispatch(const DispatcherCallPtr&, const ConnectionPtr&) = 0;
};

typedef IceUtil::Handle<Dispatcher> DispatcherPtr;
```

The Ice run time invokes the `dispatch` method whenever an operation invocation arrives or an AMI invocation completes, passing an instance of `DispatcherCall` and the connection via which the invocation arrived. The job of `dispatch` is to pass the incoming invocation to an operation implementation.

The connection parameter allows you to decide how to dispatch the operation based on the connection via which it was received. This value may be nil if no connection currently exists.

You can write `dispatch` such that it blocks and waits for completion of the invocation because `dispatch` is called by a thread in the server-side thread pool (for incoming operation invocations) or the client-side thread pool (for AMI callbacks).

The `DispatcherCall` instance encapsulates all the details of the incoming call. It is another abstract base class with the following interface:

C++

```
class DispatcherCall : virtual public IceUtil::Shared
{
public:
    virtual ~DispatcherCall() { }

    virtual void run() = 0;
};

typedef IceUtil::Handle<DispatcherCall> DispatcherCallPtr;
```

Your implementation of `dispatch` is expected to call `run` on the `DispatcherCall` instance (or, more commonly, to cause `run` to be called some time later). When you call `run`, the Ice run time processes the invocation in the thread that calls `run`.

A very simple way to implement `dispatch` would be as follows:

Copyright 2016, ZeroC, Inc.

C++

```
class MyDispatcher : public Ice::Dispatcher
public:
    virtual void dispatch(const Ice::DispatcherCallPtr& d,
const Ice::ConnectionPtr&)
    {
        d->run(); // Does not throw, blocks until op completes.
    }
};
```

Whenever the Ice run time receives an incoming operation invocation or when an AMI invocation completes, it calls `dispatch` which, in turn, calls `run` on the `DispatcherCall` instance.

With this simple example, `dispatch` immediately calls `run`, and `run` does not return until the corresponding operation invocation is complete. As a result, this implementation ties up a thread in the thread pool for the duration of the call.

So far, we really have not gained anything because all we have is a callback method that is called by the Ice run time. However, this simple mechanism is sufficient to ensure that we can update a UI from the correct thread.

A common technique to avoid blocking is to use asynchronous method invocation. In response to a UI event (such as the user pressing a "Submit" button), the application initiates an operation invocation from the corresponding UI callback by calling the operation's `begin_` meth od. This is guaranteed not to block the caller, so the UI remains responsive. Some time later, when the operation completes, the Ice run time invokes an AMI callback from one of the threads in its thread pool. That callback now has to update the UI, but that can only be done from the UI thread. By using a dispatcher, you can easily delegate the update to the correct thread. For example, here is how you can arrange for AMI callbacks to be passed to the UI thread with MFC:

C++

```
class MyDialog : public CDialog { ... };

class MyDispatcher : public Ice::Dispatcher {
public:
    MyDispatcher(MyDialog* dialog) : _dialog(dialog)
    {
    }

    virtual void
    dispatch(const Ice::DispatcherCallPtr& call, const Ice::ConnectionPtr&)
    {
        _dialog->PostMessage(WM_AMI_CALLBACK, 0,

reinterpret_cast<LPARAM>(new Ice::DispatcherCallPtr(call)));
    }

private:
    MyDialog* _dialog;
};
```

The `MyDispatcher` class simply stores the `CDialog` handle for the UI and calls `PostMessage`, passing the `DispatcherCall` instance. In turn, this causes the UI thread to receive an event and invoke the UI callback method that was registered to respond to `WM_AMI_CALLBAC K` events.

Copyright 2016, ZeroC, Inc.

In turn, the implementation of the callback method calls `run`:

C++

```
LRESULT
MyDialog::OnAMICallback(WPARAM, LPARAM lParam)
{
    try {
        Ice::DispatcherCallPtr* call =
reinterpret_cast<Ice::DispatcherCallPtr*>(lParam);
        (*call)->run();
        delete call;
    } catch (const Ice::Exception& ex) {
        // ...
    }
    return 0;
}
```

The Ice run time calls `dispatch` once the asynchronous operation invocation is complete. In turn, this causes the `OnAMICallback` to trigger, which calls `run`. Because the operation has completed already, `run` does not block, so the UI remains responsive.

Please see the `MFC` demo in your Ice distribution for a fully-functional UI client that uses this technique.

Dispatching with C++11 Lambda Functions

With a suitable C++ compiler and an Ice installation that was built with C++11 features enabled, you can use lambda functions instead of defining a callback class. The dispatcher API includes the following helper function:

C++

```
namespace Ice {
inline DispatcherPtr
newDispatcher(const std::function<void (const Ice::DispatcherCallPtr&, const
Ice::ConnectionPtr)>& cb)
{
    return ...;
}
}
```

The return value of `newDispatcher` is a `Dispatcher` object that delegates to your lambda function. You can invoke `newDispatcher` as shown below:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::DispatcherPtr disp = Ice::newDispatcher([](const Ice::DispatcherCallPtr& d,
                                                const Ice::ConnectionPtr&) {
    d->run();
});
initData.dispatcher = disp;
communicator = Ice::initialize(...);
```

Java Dispatcher API

To install a dispatcher, you must instantiate a class that implements `Ice.Dispatcher` and initialize a communicator with that instance in the `InitializationData` structure. All invocations that arrive for this communicator are made via the specified dispatcher. For example:

Java

```
public class MyDispatcher implements Ice.Dispatcher
{
    // ...
}

public class Server
{
    public static void
    main(String[] args)
    {
        Ice.Communicator communicator;

        try {
            Ice.InitializationData initData = new Ice.InitializationData();
            initData.properties = Ice.Util.createProperties(args);
            initData.dispatcher = new MyDispatcher();
            communicator = Ice.Util.initialize(args, initData);

            // ...
        } catch (Ice.LocalException & ex) { {
            // ...
        }

        // ...
    }

    // ...
}
```

The `Ice.Dispatcher` interface looks as follows:

Copyright 2016, ZeroC, Inc.

Java

```
public interface Dispatcher
{
    void dispatch(Runnable runnable, Ice.Connection con);
}
```

The Ice run time invokes the `dispatch` method whenever an operation invocation arrives, passing a `Runnable` and the connection via which the invocation arrived. The job of `dispatch` is to pass the incoming invocation to an operation implementation.

The connection parameter allows you to decide how to dispatch the operation based on the connection via which it was received. This value may be null if no connection currently exists.

You can write `dispatch` such that it blocks and waits for completion of the invocation because `dispatch` is called by a thread in the server-side thread pool (for incoming operation invocations) or the client-side thread pool (for AMI callbacks).

Your implementation of `dispatch` is expected to call `run` on the `Runnable` instance (or, more commonly, to cause `run` to be called some time later). When you call `run`, the Ice run time processes the invocation in the thread that calls `run`.

A very simple way to implement `dispatch` would be as follows:

Java

```
public class MyDispatcher implements Ice.Dispatcher
{
    public void
    dispatch(Runnable runnable, Ice.Connection connection)
    {
        // Does not throw, blocks until op completes.
        runnable.run();
    }
}
```

Whenever the Ice run time receives an incoming operation invocation or when an AMI invocation completes, it calls `dispatch` which, in turn, calls `run` on the `Runnable` instance.

With this simple example, `dispatch` immediately calls `run`, and `run` does not return until the corresponding operation invocation is complete. As a result, this implementation ties up a thread in the thread pool for the duration of the call.

So far, we really have not gained anything because all we have is a callback method that is called by the Ice run time. However, this simple mechanism is sufficient to ensure that we can update a UI from the correct thread.

A common technique to avoid blocking is to use asynchronous method invocation. In response to a UI event (such as the user pressing a "Submit" button), the application initiates an operation invocation from the corresponding UI callback by calling the operation's `begin_` meth od. This is guaranteed not to block the caller, so the UI remains responsive. Some time later, when the operation completes, the Ice run time invokes an AMI `response` callback from one of the threads in its thread pool. That callback now has to update the UI, but that can only be done from the UI thread. By using a dispatcher, you can easily delegate the update to the correct thread. For example, here is how you can arrange for AMI callbacks to be passed to the UI thread with Swing:

Copyright 2016, ZeroC, Inc.

Java

```
public class Client extends JFrame
{
    public static void main(final String[] args)
    {
        SwingUtilities.invokeLater(new Runnable()
        {
            public void run()
            {
                try {
                    new Client(args);
                } catch (Ice.LocalException e) {
                    JOptionPane.showMessageDialog(
                        null, e.toString(),
                        "Initialization failed",
                        JOptionPane.ERROR_MESSAGE);
                }
            }
        });
    }

    Client(String[] args)
    {
        Ice.Communicator communicator;

        try {
            Ice.InitializationData initData = new Ice.InitializationData();
            initData.dispatcher = new Ice.Dispatcher()
            {
                public void
                dispatch(Runnable runnable, Ice.Connection connection)
                {
                    SwingUtilities.invokeLater(runnable);
                }
            };
            communicator = Ice.Util.initialize(args, initData);
        }
        catch(Throwable ex)
        {
            // ...
        }
        // ...
    }

    // ...
}
```

The `dispatch` method simply delays the call to `run` by calling `invokeLater`, passing it the `Runnable` that is provided by the Ice run time. This causes the Swing UI thread to eventually make the call to `run`. Because the Ice run time does not call `dispatch` until the asynchronous invocation is complete, that call to `run` does not block and the UI remains responsive.

Please see the `swing` demo in your Ice distribution for a fully-functional UI client that uses this technique.

Copyright 2016, ZeroC, Inc.

C# Dispatcher API

To install a dispatcher, you must initialize a communicator with a delegate of type `Ice.Dispatcher` in the `InitializationData` structure. All invocations that arrive for this communicator are made via the specified dispatcher. For example:

C#

```
public class Server
{
    public static void Main(string[] args)
    {
        Ice.Communicator communicator = null;

        try {
            Ice.InitializationData initData = new Ice.InitializationData();
            initData.dispatcher = new MyDispatcher().dispatch;
            communicator = Ice.Util.initialize(ref args, initData);
            // ...
        } catch (System.Exception ex) {
            // ...
        }

        // ...
    }

    // ...
}
```

The `Ice.Dispatcher` delegate is defined as follows:

C#

```
public delegate void Dispatcher(System.Action call, Connection con);
```

The Ice run time calls your delegate whenever an operation invocation arrives, passing a `System.Action` delegate and the connection via which the invocation arrived. The job of your delegate is to pass the incoming invocation to an operation implementation.

The connection parameter allows you to decide how to dispatch the operation based on the connection via which it was received. This value may be null if no connection currently exists.

In this example, the delegate calls a method `dispatch` on an instance of a `MyDispatcher` class. You can write `dispatch` such that it blocks and waits for completion of the invocation because `dispatch` is called by a thread in the server-side thread pool (for incoming operation invocations) or the client-side thread pool (for AMI callbacks).

Your implementation of `dispatch` is expected to invoke the `call` delegate (or, more commonly, to cause it to be invoked some time later). When you invoke the `call` delegate, the Ice run time processes the invocation in the thread that invokes the delegate.

A very simple way to implement `dispatch` would be as follows:

Copyright 2016, ZeroC, Inc.

C#

```
public class MyDispatcher
{
    public void
    dispatch(System.Action call, Ice.Connection con)
    {
        // Does not throw, blocks until op completes.
        call();
    }
};
```

Whenever the Ice run time receives an incoming operation invocation or when an AMI invocation completes, it calls `dispatch` which, in turn, invokes the `call` delegate.

With this simple example, `dispatch` immediately invokes the delegate, and that call does not return until the corresponding operation invocation is complete. As a result, this implementation ties up a thread in the thread pool for the duration of the call.

So far, we really have not gained anything because all we have is a callback method that is called by the Ice run time. However, this simple mechanism is sufficient to ensure that we can update a UI from the correct thread.

A common technique to avoid blocking is to use asynchronous method invocation. In response to a UI event (such as the user pressing a "Submit" button), the application initiates an operation invocation from the corresponding UI callback by calling the operation's `begin_` meth od. This is guaranteed not to block the caller, so the UI remains responsive. Some time later, when the operation completes, the Ice run time invokes an AMI callback from one of the threads in its thread pool. That callback now has to update the UI, but that can only be done from the UI thread. By using a dispatcher, you can easily delegate the update to the correct thread. For example, here is how you can arrange for AMI callbacks to be passed to the UI thread with WPF:

Copyright 2016, ZeroC, Inc.

C#

```
public partial class MyWindow : Window
{
    private void Window_Loaded(object sender, EventArgs e)
    {
        Ice.Communicator communicator = null;

        try
        {
            Ice.InitializationData initData = new Ice.InitializationData();
            initData.dispatcher =
                    delegate(System.Action action, Ice.Connection connection)
            {
                Dispatcher.BeginInvoke(DispatcherPriority.Normal, action);
            };
            communicator = Ice.Util.initialize(initData);
        }
        catch(Ice.LocalException ex)
        {
            // ...
        }
    }

    // ...
}
```

The delegate calls `Dispatcher.BeginInvoke` on the `action` delegate. This causes WPF to queue the actual asynchronous invocation of `action` for later execution by the UI thread. Because the Ice run time does not invoke your delegate until an asynchronous operation invocation is complete, when the UI thread executes the corresponding call to the `EndInvoke` method, that call does not block and the UI remains responsive.

The net effect is that you can invoke an operation asynchronously from a UI callback method without the risk of blocking the UI thread. For example:

Copyright 2016, ZeroC, Inc.

C#

```
public partical class MyWindow : Window
{
    private void someOp_Click(object sender, RoutedEventArgs e)
    {
        MyIntfPrx p = ...;

        // Call remote operation asynchronously.
        // Response is processed in UI thread.
        p.begin_someOp().whenCompleted(this.opResponse, this.opException);
    }

    public void opResponse()
    {
        // Update UI...
    }

    public void opException(Ice.Exception ex)
    {
        // Update UI...
    }
}
```

Please see the `wpf` demo in your Ice distribution for a fully-functional UI client that uses this technique.

Objective-C Dispatcher API

To install a dispatcher, you must initialize a communicator with a callback (as an Objective-C block) in the `ICEInitializationData` structure. All invocations that arrive for this communicator are made via the specified callback. For example:

Copyright 2016, ZeroC, Inc.

Objective-C

```
int
main(int argc, char* argv[])
{
    objc_startCollectorThread();
    id<ICECommunicator> communicator = nil;
    @try
    {
        ICEInitializationData* initData = [ICEInitializationData
initializationData];
        initData.dispatcher =
            ^(id<ICEDispatcherCall> call, id<ICEConnection> con)
            {
                // ...
            };
        communicator = [ICEUtil createCommunicator:&argc argv:argv
initData:initData];
        // ...
    }
    @catch(ICELocalException* ex)
    {
        // ...
    }

    // ...
}
```

The type of the dispatcher callback must match the following block signature:

Objective-C

```
void(^)(id<ICEDispatcherCall> call, id<ICEConnection> connection)
```

The Ice run time invokes the dispatcher callback whenever an operation invocation arrives, passing an object implementing the `ICEDispatcherCall` protocol and the connection via which the invocation arrived. The job of your callback implementation is to execute the given call.

The connection parameter allows you to decide how to dispatch the operation based on the connection via which it was received. This value may be nil if no connection currently exists.

You can write the callback such that it blocks and waits for completion of the invocation because the callback is called by a thread in the server-side thread pool (for incoming operation invocations) or the client-side thread pool (for AMI callbacks).

The `ICEDispatcherCall` protocol defines how to execute the incoming call:

Copyright 2016, ZeroC, Inc.

Objective-C

```
@protocol ICEDispatcherCall <NSObject>
-(void) run;
@end
```

Your callback is expected to call `run` on the `ICEDispatcherCall` instance (or, more commonly, to cause run to be called some time later). When you call `run`, the Ice run time processes the invocation in the thread that calls `run`.

A very simple way to implement the dispatcher callback would be as follows:

Objective-C

```
void(^myDispatcher)(id<ICEDispatcherCall>, id<ICEConnection>) =
    ^(id<ICEDispatcherCall> call, id<ICEConnection> con)
    {
        // Does not throw, blocks until op completes.
        [call run];
    };
```

Whenever the Ice run time receives an incoming operation invocation or when an AMI invocation completes, it calls the dispatcher callback which, in turn, invokes the run method on the call.

With this simple example, the dispatcher callback immediately invokes `run`, and `run` does not return until the corresponding operation invocation is complete. As a result, this implementation ties up a thread in the thread pool for the duration of the call.

So far, we really have not gained anything because all we have is a callback method that is called by the Ice run time. However, this simple mechanism is sufficient to ensure that we can update a UI from the correct thread.

A common technique to avoid blocking is to use asynchronous method invocation. In response to a UI event (such as the user pressing a "Submit" button), the application initiates an operation invocation from the corresponding UI callback by calling the operation's `begin_` meth od. This is guaranteed not to block the caller, so the UI remains responsive. Some time later, when the operation completes, the Ice run time invokes an AMI callback from one of the threads in its thread pool. That callback now has to update the UI, but that can only be done from the UI thread. By using a dispatcher, you can easily delegate the update to the correct thread. For example, here is how you can arrange for AMI callbacks to be passed to Cocoa or Cocoa Touch main thread:

Copyright 2016, ZeroC, Inc.

Objective-C

```
-(void)viewDidLoad
{
    ICEInitializationData* initData = [ICEInitializationData initializationData];
    initData.dispatcher =
        ^(id<ICEDispatcherCall> call, id<ICEConnection> con)
        {
            dispatch_sync(dispatch_get_main_queue(), ^ { [call run]; });
        };

    communicator = [[ICEUtil createCommunicator:initData] retain];

    // ....
}
```

The dispatcher callback calls `dispatch_sync` on the main queue. This queues the actual call for later execution by the main thread. Because the Ice run time does not invoke the dispatcher callback until an asynchronous operation invocation is complete, when the UI thread executes the corresponding call, that call does not block and the UI remains responsive.

The net effect is that you can invoke an operation asynchronously from a UI callback method without the risk of blocking the UI thread. For example:

Objective-C

```
-(void)someOp:(id)sender
{
    id<MyIntfPrx> p = ...;
    [p begin_someOp:^{ [self response]; }
          exception:^(ICEException* ex) { [self exception:ex]; }];
}

-(void) response
{
    // Update UI...
}

-(void) exception:(ICEException* ex)
{
    // Update UI...
}
```

Please see the Cocoa or iPhone demos in your Ice Touch distribution for fully-functional UI clients that use this technique.

Dispatcher Implementation Notes

An application that uses a custom dispatcher must adhere to the following rules to avoid a deadlock:

- Dispatcher implementations must ensure that all requests are dispatched. Failing to dispatch all requests will cause `Communicator::destroy` to hang indefinitely. If a dispatcher has resources that must be reclaimed (e.g., joining with a helper thread), it can safely do so after `Communicator::destroy` has completed.

Copyright 2016, ZeroC, Inc.

- Never make a blocking invocation from the dispatch thread, such as a synchronous proxy operation or a proxy method that can potentially block, such as `ice_getConnection`. These invocations depend on the dispatcher for their own completion, therefore blocking the dispatch thread will inevitably lead to deadlock.

See Also

- Asynchronous Method Invocation (AMI) in C++
- Asynchronous Method Invocation (AMI) in Java
- Asynchronous Method Invocation (AMI) in C-Sharp
- Asynchronous Method Invocation (AMI) in Objective-C

Copyright 2016, ZeroC, Inc.

Blocking API Calls

This page lists the Ice APIs that can potentially block the calling thread. Graphical applications should normally avoid calling these APIs from the "event loop" thread.

The JavaScript mapping does not have any blocking APIs.

All synchronous remote invocations can block, including the built-in proxy operations `ice_ping`, `ice_isA`, `ice_id` and `ice_ids`. Since `checkedCast` internally calls `ice_isA`, it too can block.

For asynchronous remote invocations, the `begin_` method never blocks but the `end_` method can block.

Ice also has a number of local APIs that can block, either because they wait indefinitely for a condition to become true, or because their implementations may make remote invocations:

- `AsyncResult`
  - `waitForCompleted`
  - `waitForSent`
- `Communicator`
  - `createAdmin`
  - `createObjectAdapter`
  - `createObjectAdapterWithEndpoints`
  - `createObjectAdapterWithRouter`
  - `destroy`
  - `end_flushBatchRequests`
  - `flushBatchRequests`
  - `getAdmin`
  - `waitForShutdown`
- `Connection`
  - `close`
  - `end_flushBatchRequests`
  - `flushBatchRequests`
  - `waitUntilHolding`
- `ObjectPrx`
  - `ice_getConnection`
- `ObjectAdapter`
  - `activate`
  - `deactivate`
  - `destroy`
  - `waitForDeactivate`
  - `waitForHold`

The Java mapping allows you to interrupt threads blocked in these APIs.

See Also

- Java Interrupts

Copyright 2016, ZeroC, Inc.

Connection Management

The Ice run time establishes connections automatically and transparently as a side effect of using proxies. There are well-defined rules that determine when a new connection is established. If necessary, you can influence connection management activities.

Connection management becomes increasingly important as network environments grow more complex. In particular, if you need to make callbacks from a server to a client through a firewall, you must use a bidirectional connection. In most cases, you can use a Glacier2 router t o automatically take advantage of bidirectional connections. However, the Ice run time also provides direct access to connections, allowing you to explicitly control establishment and closure of both unidirectional and bidirectional connections.

Most Ice applications benefit from active connection management and transparent connection establishment and thus need not concern themselves with the details of connections. Not all Ice applications can be so fortunate, and for those applications Ice provides convenient ac cess to connections that enables developers to address the realities of today's deployment environments.

The discussion that follows assumes that you are familiar with proxies and endpoints.

Topics

- Connection Establishment
- Active Connection Management
- Using Connections
- Connection Closure
- Bidirectional Connections

Copyright 2016, ZeroC, Inc.

Connection Establishment

Connections are established as a side effect of using proxies. The first invocation on a proxy causes the Ice run time to search for an existing connection to one of the proxy's endpoints; only if no suitable connection exists does the Ice run time establish a new connection to one of the proxy's endpoints.

This page describes how and when Ice establishes a new connection.

On this page:

- Endpoint Selection for New Connections
- Interface Selection for New Connections
- Error Semantics for Failed Connections
- Reusing an Existing Connection
  - Connection Reuse for Proxies with Multiple Endpoints
  - Protocol Compression and Connection Reuse
  - Influencing Connection Reuse
- Connection Caching
- Timeouts and Connection Establishment

Endpoint Selection for New Connections

A proxy performs a number of operations on its endpoints before it asks the Ice run time to supply a connection. These operations produce a list of zero or more endpoints that satisfy the proxy's configuration. If the resulting list is empty, the application receives `NoEndpointException` to indicate that no suitable endpoints could be found. For example, this situation can arise when a twoway proxy contains only a UDP endpoint; the UDP endpoint is eliminated from consideration because it cannot be used for twoway invocations.

The proxy performs the following steps to derive its endpoint list:

1. Remove the endpoints of unknown transports. For instance, SSL endpoints are removed if the SSL plug-in is not installed.
2. Remove endpoints that are not suitable for the proxy's invocation mode. For example, datagram endpoints are removed for twoway, oneway and batch oneway proxies. Similarly, non-datagram endpoints are removed for datagram and batch datagram proxies.
3. Perform DNS queries to convert host names into IP addresses, if necessary. For a multi-homed host name, the proxy adds a new endpoint for each address returned by the DNS query.
4. Sort the endpoints according to the configured selection type, which is established using the `ice_endpointSelection` proxy method. The default value is `Random`, meaning the endpoints are randomly shuffled. Alternatively, the value `Ordered` maintains the existing order of the endpoints.
5. Satisfy the proxy's security requirements:
   - If `Ice.Override.Secure` is defined, remove all non-secure endpoints.
   - Otherwise, if the proxy is configured to prefer secure endpoints (e.g., by calling the `ice_preferSecure` proxy method), move all secure endpoints to the beginning of the list. Note that this setting still allows non-secure endpoints to be included.
   - Otherwise, move all non-secure endpoints to the beginning of the list.

If connection caching is enabled and the Ice run time already has a compatible connection, it reuses the cached connection. Otherwise, the run time attempts to connect to each endpoint in the list until it succeeds or exhausts the list; the order in which endpoints are selected for connection attempts depends on the endpoint selection policy. This policy can be set using a default property (`Ice.Default.EndpointSelection`), using a proxy property (*`name`*`.EndpointSelection`), and using the `ice_endpointSelection` proxy method.

Interface Selection for New Connections

If a new outgoing connection must be established, Ice's default behavior allows the operating system to select the network interface to be used for the connection. You can force Ice to use a specific network interface for outgoing connections by configuring a *source address*. There are two ways to do this:

- Define the `Ice.Default.SourceAddress` property to establish a default source address for all outgoing connections created by a communicator
- Include a `--sourceAddress` option in proxy endpoints, which overrides any setting for `Ice.Default.SourceAddress`

The value in each case must be an IP address associated with the desired network interface.

Error Semantics for Failed Connections

If a failure occurs during a connection attempt, the Ice run time tries to connect to all of the proxy's remaining endpoints until either a connection is successfully established or all attempts have failed. At that point, the Ice run time may attempt automatic retries depending on the value of the `Ice.RetryIntervals` configuration property. The default value of this property is `0`, which causes the Ice run time to try

Copyright 2016, ZeroC, Inc.

connecting to all of the endpoints one more time.

Tip
Define the property `Ice.Trace.Retry=2` to monitor these attempts.

If no connection can be established on this second attempt, the Ice run time raises an exception that indicates the reason for the final failed attempt (typically `ConnectFailedException`). Similarly, if a connection was lost during a request and could not be reestablished (assuming the request can be retried), the Ice run time raises an exception that indicates the reason for the final failed attempt.

Reusing an Existing Connection

When establishing a connection for a proxy, the Ice run time reuses an existing connection under the following conditions:

- The remote endpoint matches one of the proxy's endpoints.
- The connection was established by the communicator that created the proxy.
- The connection matches the proxy's configuration. Timeout values play an important role here, as an existing connection is only reused if its timeout value (i.e., the timeout used when the connection was established) matches the endpoint timeout in the new proxy. Similarly, a proxy configured with a connection ID only reuses a connection if it was established by a proxy with the same connection ID.

Connection Reuse for Proxies with Multiple Endpoints

Applications must exercise caution when using proxies containing multiple endpoints, especially endpoints using different transports. For example, suppose a proxy has multiple endpoints, such as one each for TCP, SSL, and UDP. When establishing a connection for this proxy, the Ice run time will open a new connection only if it cannot reuse an existing connection to any of the endpoints (that is, if connection caching is enabled). Furthermore, the proxy in its default (that is, non-secure) configuration gives higher priority to non-secure endpoints. If you want to ensure that a particular transport is used by a proxy, you must configure the proxy appropriately, such as by calling the proxy methods `ice_secure` or `ice_datagram` as necessary.

Protocol Compression and Connection Reuse

The Ice run time does not consider compression settings when searching for existing connections to reuse; proxies whose compression settings differ can share the same connection (assuming all other selection criteria are satisfied).

Influencing Connection Reuse

The default behavior of the Ice run time, which reuses connections whenever possible, is appropriate for many applications because it conserves resources and typically has little or no impact on performance. However, when a server implementation attaches semantics to a connection, the client often must be designed to cooperate, despite the tighter coupling it causes. For example, a server might use a serialized thread pool to preserve the order of requests received over each connection. If the client wants to execute several requests simultaneously, it must be able to force the Ice run time to establish new connections at will.

For those situations that require more control over connection reuse, the Ice run time allows you to form arbitrary groups of proxies that share a connection by configuring them with the same connection identifier. The proxy method `ice_connectionId` returns a new proxy configured with the given connection ID. Once configured, the Ice run time ensures that the proxy only reuses a connection that was established by a proxy with the same connection ID (assuming all other criteria for connection reuse are also satisfied). A new connection is created if none with a matching ID is found, which means each proxy could conceivably have its own connection if each were assigned a unique connection ID.

As an example, consider the following code fragment:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx prx = comm->stringToProxy("ident:tcp -p 10000");
Ice::ObjectPrx g1 = prx->ice_connectionId("group1");
Ice::ObjectPrx g2 = prx->ice_connectionId("group2");
prx->ice_ping(); // Opens a new connection
g1->ice_ping(); // Opens a new connection
g2->ice_ping(); // Opens a new connection
MyInterfacePrx i1 = MyInterfacePrx::checkedCast(g1);
i1->ice_ping(); // Reuses g1's connection
MyInterfacePrx i2 = MyInterfacePrx::checkedCast(prx->ice_connectionId("group2"));
i2->ice_ping(); // Reuses g2's connection
```

A total of three connections are established by this example:

1. The proxy `prx` establishes a new connection. This proxy has the default connection ID (an empty string).
2. The proxy `g1` establishes a new connection because the only existing connection, the one established by `prx`, has a different connection ID.
3. Similarly, the proxy `g2` establishes a new connection because none of the existing connections have a matching connection ID.

The proxy `i1` inherits its connection ID from `g1`, and therefore shares the connection for `group1`; `i2` explicitly configured its connection ID and shares the `group2` connection with proxy `g2`.

Connection Caching

When we refer to a proxy's connection, we actually mean the connection that the proxy is *currently* using. This connection can change over time, such that a proxy might use several connections during its lifetime. For example, an idle connection may be closed automatically and then transparently replaced by a new connection when activity resumes.

After establishing a connection in response to proxy activities, the Ice run time adds the connection to an internal pool for subsequent reuse by other proxies. The Ice run time manages the lifetime of the connection and eventually closes it. The connection is not affected by the life cycle of the proxies that use it, except that the lack of activity may prompt the Ice run time to close the connection after a while.

Once a proxy has been associated with a connection, the proxy's default behavior is to continue using that connection for all subsequent requests. In effect, the proxy caches the connection and attempts to use it for as long as possible in order to minimize the overhead of creating new connections. If the connection is later closed and the proxy is used again, the proxy repeats the connection-establishment procedure described earlier.

There are situations in which this default caching behavior is undesirable, such as when a client has a proxy with multiple endpoints and wishes to balance the load among the servers at those endpoints. The client can disable connection caching by passing an argument of `false` to the proxy method `ice_connectionCached`. The new proxy returned by this method repeats the connection-establishment procedure before each request, thereby achieving request load balancing at the expense of potentially higher latency.

This type of load balancing is performed solely by the client using whatever endpoints are contained in the proxy. More sophisticated forms of load balancing are also possible, such as when using IceGrid.

Timeouts and Connection Establishment

As of Ice 3.6, the default timeout for all connections is 60 seconds, as determined by the `Ice.Default.Timeout` property. The connection timeout value applies to all network operations. If a connection cannot be established within the allotted time, Ice raises `ConnectTimeoutException`.

You can set a timeout on a proxy using the `ice_timeout` proxy method. To use the same timeout period for all proxies, you can define the `Ice.Override.Timeout` property; in this case, Ice ignores any timeout established using the `ice_timeout` proxy method or the `Ice.Default.Timeout` property. Finally, if you want to specify a separate timeout value that affects only connection establishment and takes precedence over a proxy's configured timeout value, you can define the `Ice.Override.ConnectTimeout` property.

Copyright 2016, ZeroC, Inc.

Connection timeout values affect connection reuse. For example, if the endpoint in proxy A is identical to the endpoint in proxy B except their timeout values differ, the proxies cannot share the same connection.

The timeout in effect when a connection is established is bound to that connection and cannot be changed. If a network operation times out, all outstanding requests on that connection receive a `TimeoutException` and the connection is closed forcefully. The Ice run time automatically retries these requests on a new connection, assuming that automatic retries are enabled and would not violate at-most-once semantics.

Invocation timeouts are a separate feature added in Ice 3.6 and do not affect connection reuse.

See Also

- Proxy Methods
- Proxy Endpoints
- The Ice Threading Model
- Automatic Retries
- Connection Timeouts
- Active Connection Management
- IceGrid
- Ice.Default.*
- Proxy Properties
- Miscellaneous Ice.* Properties

Copyright 2016, ZeroC, Inc.

Active Connection Management

Active Connection Management (ACM) is enabled by default and helps to improve scalability and conserve application resources by closing idle connections.

On this page:

- Configuring Active Connection Management
- ACM Timeout Semantics
- Creating a Connection Management Policy
  - Bidirectional Connections
  - Managing Sessions
  - Detecting Peer Problems
  - Disabling ACM

Configuring Active Connection Management

There are three components to an ACM configuration:

- Close
  Determines the situations in which Ice closes a connection
- Heartbeat
  Determines the situations in which Ice sends "heartbeat" messages to keep a connection alive
- Timeout
  Defines the time interval in which the Close and Heartbeat actions occur

You can configure these components using ACM properties that affect all of the connections created by a communicator or modify a single connection directly in your code:

- `Ice.ACM.Close`
  `Ice.ACM.Heartbeat`
  `Ice.ACM.Timeout`
  These properties establish the default settings for a communicator and affect both client (outgoing) and server (incoming) connections.
- `Ice.ACM.Client.Close`
  `Ice.ACM.Client.Heartbeat`
  `Ice.ACM.Client.Timeout`
  These properties override the default properties above for client (outgoing) connections.
- `Ice.ACM.Server.Close`
  `Ice.ACM.Server.Heartbeat`
  `Ice.ACM.Server.Timeout`
  These properties override the default properties above for server (incoming) connections.
- *`adapter`*`.ACM.Close`
  `adapter``.ACM.Heartbeat`
  `adapter``.ACM.Timeout`
  These properties override the `Ice.ACM.Server` properties above for connections to a particular object adapter.
- Programmatically by calling `setACM` on a connection object. These settings override all ACM property configurations for that connection.

ACM Timeout Semantics

That the Close and Heartbeat behaviors share a single Timeout setting may seem surprising at first glance. For instance, it's possible for a program to configure certain combinations of Close and Heartbeat settings such that the heartbeats prevent a connection from ever becoming eligible for closure. In general, however, one peer will configure Close and Timeout settings while the other peer will configure Heartbeat and Timeout settings, in a coordinated effort to meet the application's requirements for connection management.

For each connection configured for ACM, Ice checks its status approximately every (Timeout / 2) seconds. If heartbeats are configured for the connection, Ice also sends a heartbeat during every status check. If two peers use the same Timeout value, this strategy ensures that heartbeat messages are sent well before the connection would be considered idle by the other peer.

Creating a Connection Management Policy

Copyright 2016, ZeroC, Inc.

The ACM settings give you a lot of flexibility for developing a connection management policy that meets the needs of your application. As application requirements can vary greatly, we provide recommendations for several common use cases in the subsections below. You should also take the following general guidelines into consideration when developing your policy:

- Connection semantics
  Is your application dependent on a connection remaining open? For example, there may be semantics attached to a connection, as in the case of a session in which a peer's identity, or some allocated resources, are valid only as long as the peer's connection remains active. Glacier2 sessions work this way. In this case you'd want to enable heartbeats and configure the timeout so that it is compatible with the session's expiration timeout.
- Network traversal issues
  Depending on your network topology, security requirements, and other application considerations, a connection may be a valuable resource that is expensive or time-consuming to construct and you may not want Ice to discard it lightly. If sending heartbeats is too costly, you may need to disable automatic closure altogether.
- Aggressiveness
  ACM's combination of heartbeats and connection closure options allows your application to detect when something has gone wrong with a peer. Furthermore, the connection closure options provide varying levels of response, from conservative to aggressive. Your application requirements will dictate the necessary behavior.
- Local networks
  When operating in a purely local network in which there are no firewalls to traverse and none of the other considerations listed above are a concern, you can usually allow Ice to open and close connections transparently. In fact, the default configuration is a reasonable starting point.
- Two sides to every connection
  It's important to consider the requirements of both clients *and* servers when designing your connection management policy. For example, if a client and server are using significantly different settings for their ACM timeouts, it can result in surprising and usually undesirable behavior.

To verify that the policy you've configured is behaving as expected, set `Ice.Trace.Network=2` and monitor your log output for connection-related activity.

Bidirectional Connections

A bidirectional connection allows a server to make callback invocations on a client, offering a simple solution to work around the limitations enforced by network firewalls in which the server would otherwise be prevented from establishing a separate connection back to the client. Given that the client's connection to the server represents the server's only path to that client, this connection must remain open as long as the client needs it.

In versions prior to Ice 3.6, our recommendation was to disable ACM completely in both client and server when using bidirectional connections to prevent unintended closure. As of Ice 3.6, there's no harm in enabling ACM connection closure as long as you also enable client-side heartbeats to prevent the connection from becoming idle during periods of inactivity. Consider the following settings:

```
# Client
Ice.ACM.Close=0      # CloseOff
Ice.ACM.Heartbeat=3  # HeartbeatAlways
Ice.ACM.Timeout=30

# Server
Ice.ACM.Close=4      # CloseOnIdleForceful
Ice.ACM.Heartbeat=0  # HeartbeatOff
Ice.ACM.Timeout=30
```

Here the client will always send heartbeats at regular intervals to keep connections alive. There's no need for the server to send heartbeats because the client has disabled automatic connection closure. If a connection does become idle, which most likely would be due to a serious problem with the client, the server forcefully closes the connection regardless of whether any incoming or outgoing requests are pending at the time.

If your application manually configures bidirectional connections (as opposed to the automatic setup provided by Glacier2 connections, for instance), it's not much extra work to configure the ACM settings individually on the connection object if your requirements for bidirectional connections differ from the communicator-wide settings defined by the ACM properties. As an example, an application may not always need

Copyright 2016, ZeroC, Inc.

to use a connection for bidirectional purposes. As long as a connection is unidirectional, the application might consider it safe to be closed automatically by ACM. Once it transitions to a bidirectional connection, the program that initiated the connection should modify its ACM configuration similar to the client settings shown above.

Managing Sessions

The notion of a session is a common and very useful solution for managing resources associated with a particular client. Briefly, the idea is that a client creates a session, which usually includes an authentication process, and then allocates some resources. The server associates those resources with the client's session and requires the client to keep the session active, using an expiration timer to reclaim a session and its resources if a client abandons the session without formally terminating it. The strategy represents good defensive programming for a server; without such a solution, ill-behaved clients could continue allocating resources indefinitely.

In versions prior to Ice 3.6, we recommended that the client create a background thread that periodically "pinged" the server using an interval based on the server's session expiration time. The thread is no longer necessary as of Ice 3.6, since the ACM heartbeat functionality serves the same purpose. Note however that there are still a couple of considerations:

- The client needs to determine the server's session expiration time. An application might configure this statically, or the value may only be available at run time by calling an Ice operation. In the latter case, the client would need to transfer this value to a corresponding ACM timeout setting, most likely by calling `setACM` on the connection object.
- Using ACM heartbeats to keep a session alive is convenient for clients but presents a problem for server implementations: How does a server know that the connection is still alive? The solution is for the server to call `setCallback` on each connection to register a callback that Ice invokes each time a heartbeat is received, and when the connection is eventually closed.

The following configuration settings should get you started:

```
# Client
Ice.ACM.Close=0      # CloseOff
Ice.ACM.Heartbeat=3  # HeartbeatAlways
Ice.ACM.Timeout=30

# Server
Ice.ACM.Close=4      # CloseOnIdleForceful
Ice.ACM.Heartbeat=0  # HeartbeatOff
Ice.ACM.Timeout=30
```

These settings assume that the application configures the ACM timeout statically. (For example, perhaps the ACM timeout also serves directly as the session expiration timeout.) The server's setting for ACM connection closure represents a design decision that assists the implementation: When Ice detects that a connection has become idle, likely due to a serious problem with the client, it forcefully closes the connection. This process involves calling the callback that the server previously set on the connection; presumably the server interprets this notification as an indication that the session can be destroyed. Without the connection callback, the server would have to implement some other strategy for periodically reaping expired sessions.

Detecting Peer Problems

ACM is an effective tool for detecting and recovering from catastrophic problems with a peer. Let's suppose that it's safe to enable ACM connection closure in an application for both the client and the server. Furthermore, in this application, the server can take a significant amount of time to dispatch a client invocation. The client is willing to wait as long as it takes for the invocation to complete, however, as a defensive measure, the client also wants the ability to recover in case the server becomes unresponsive. Consider these settings:

Copyright 2016, ZeroC, Inc.

```
# Client
Ice.ACM.Close=3      # CloseOnInvocationAndIdle
Ice.ACM.Heartbeat=0  # HeartbeatOff
Ice.ACM.Timeout=30

# Server
Ice.ACM.Close=1      # CloseOnIdle
Ice.ACM.Heartbeat=1  # HeartbeatOnInvocation
Ice.ACM.Timeout=30
```

The server configuration causes it to close connections when they become idle. Since the client doesn't send heartbeats, this means the client hasn't actively used the connection during the timeout period. The server configuration also sends heartbeats to the client, but only while the server is dispatching invocations. As a result, no matter how long it takes to dispatch the invocations, the Ice run time in the server will continue sending heartbeat messages as an indicator to the client that the server is still "healthy". If the Ice run time in the client determines that a connection has become idle, it either means the connection is no longer being used, or if outgoing invocations are still pending, it means that the server has stopped sending heartbeats for some reason. In the former case, the connection is transparently closed without affecting the client. In the latter case, Ice forcefully closes the connection. Assuming the subsequent automatic retry behavior fails, all pending client invocations on that connection will terminate with an exception.

Don't use ACM timeouts as a mechanism for aborting long-running invocations. Ice provides invocation timeouts for that purpose.

Disabling ACM

An application may be interested in preventing ACM from closing connections. For example, oneway requests can be silently discarded when a server closes a connection. One solution is to set the following property in the server:

```
Ice.ACM.Server.Close=0
```

This setting prevents the server from closing an incoming connection regardless of how long the connection has been idle.

Another solution is to enable heartbeats in the client so that the connection remains idle and never becomes eligible for closure while the client process is operating normally. If the client should terminate unexpectedly, the server will still be able to close the connection within a reasonable amount of time without waiting for a low-level notification from the network layer.

See Also

- Ice.ACM.*
- Oneway Invocations
- Object Adapters
- Batched Invocations
- Using Connections
- Connection Closure
- Bidirectional Connections
- Glacier2

Copyright 2016, ZeroC, Inc.

Using Connections

Applications can gain access to an Ice object representing an established connection.

On this page:

- The Connection Interface
  - Flushing Batch Requests for a Connection
- The Endpoint Interface
  - Opaque Endpoints
- Client-Side Connection Usage
- Server-Side Connection Usage
- Closing a Connection
  - Graceful Closure
  - Forceful Closure

The `Connection` Interface

The Slice definition of the `Connection` interface is shown below:

Slice

```
module Ice {
    local class ConnectionInfo {
        bool incoming;
        string adapterName;
        string connectionId;
    };

    local interface ConnectionCallback {
        void heartbeat(Connection con);
        void closed(Connection con);
    };

    local enum ACMClose {
        CloseOff,
        CloseOnIdle,
        CloseOnInvocation,
        CloseOnInvocationAndIdle,
        CloseOnIdleForceful
    };

    local enum ACMHeartbeat {
        HeartbeatOff,
        HeartbeatOnInvocation,
        HeartbeatOnIdle,
        HeartbeatAlways
    };

    local struct ACM {
        int timeout;
        ACMClose close;
        ACMHeartbeat heartbeat;
    };

    local interface Connection {
```


Copyright 2016, ZeroC, Inc.

```
        void close(bool force);
        Object* createProxy(Identity id);
        void setAdapter(ObjectAdapter adapter);
        ObjectAdapter getAdapter();
        Endpoint getEndpoint();
        void flushBatchRequests();
        void setCallback(ConnectionCallback callback);
        void setACM(optional(1) int timeout, optional(2) ACMClose close,
optional(3) ACMHeartbeat heartbeat);
        ACM getACM();
        string type();
        int timeout();
        string toString();
        ConnectionInfo getInfo();
    };

    local class IPConnectionInfo extends ConnectionInfo {
        string localAddress;
        int localPort;
        string remoteAddress;
        int remotePort;
    };

    local class TCPConnectionInfo extends IPConnectionInfo {};

    local class UDPConnectionInfo extends IPConnectionInfo {
        string mcastAddress;
        int mcastPort;
    };

    local class WSConnectionInfo extends IPConnectionInfo {};
};

module IceSSL {
    local class ConnectionInfo extends Ice::IPConnectionInfo {
        string cipher;
        Ice::StringSeq certs;
```


Copyright 2016, ZeroC, Inc.

```
    };
};
```

As indicated in the Slice definition, a connection is a local interface, similar to a communicator or an object adapter. A connection therefore is only usable within the process and cannot be accessed remotely.

The `Connection` interface supports the following operations:

- `void close(bool force)`
  Explicitly closes the connection. The connection is closed gracefully if `force` is false, otherwise the connection is closed forcefully.
- `Object* createProxy(Identity id)`
  Creates a special proxy that only uses this connection. This operation is primarily used for bidirectional connections.
- `void setAdapter(ObjectAdapter adapter)`
  Associates this connection with an object adapter to enable a bidirectional connection.
- `ObjectAdapter getAdapter()`
  Returns the object adapter associated with this connection, or nil if no association has been made.
- `Endpoint getEndpoint()`
  Returns an `Endpoint` object.
- `void flushBatchRequests()`
  Flushes any pending batch requests for this connection.
- `void setCallback(ConnectionCallback callback)`
  Associates a callback with this connection that is invoked whenever the connection receives a heartbeat message or is closed. Passing a nil value clears the current callback.
- `void setACM(optional(1) int timeout, optional(2) ACMClose close, optional(3) ACMHeartbeat heartbeat)`
  Configures Active Connection Management settings for this connection. All arguments are optional, therefore you can change some of the settings while leaving the others unaffected. Refer to your language mapping for more details on optional parameters.
- `ACM getACM()`
  Returns the connection's current settings for Active Connection Management.
- `string type()`
  Returns the connection type as a string, such as `"tcp"`.
- `int timeout()`
  Returns the timeout value used when the connection was established.
- `string toString()`
  Returns a readable description of the connection.
- `ConnectionInfo getInfo()`
  This operation returns a `ConnectionInfo` class defined as follows:

Slice

```
local class ConnectionInfo {
    bool incoming;
    string adapterName;
};
```

The `incoming` member is true if the connection is an incoming connection and false, otherwise. If `incoming` is true, `adapterName` provides the name of the object adapter that accepted the connection. Note that the object returned by `getInfo` implements a more derived interface, depending on the connection type. You can down-cast the returned class instance and access the connection-specific information according to the type of the connection.

Copyright 2016, ZeroC, Inc.

Flushing Batch Requests for a Connection

The `flushBatchRequests` operation blocks the calling thread until any batch requests that are queued for a connection have been successfully written to the local transport. To avoid the risk of blocking, you can also invoke this operation asynchronously using the `begin_flushBatchRequests` method (in those language mappings that support it).

Since batch requests are inherently oneway invocations, the `begin_flushBatchRequests` method does not support a request callback. However, you can use the exception callback to handle any errors that might occur while flushing, and the sent callback to receive notification that the batch request has been flushed successfully.

For example, the code below demonstrates how to flush batch requests asynchronously in C++:

C++

```
class FlushCallback : public IceUtil::Shared
{
public:

    void exception(const Ice::Exception& ex)
    {
        cout << "Flush failed: " << ex << endl;
    }

    void sent(bool sentSynchronously)
    {
        cout << "Batch sent!" << endl;
    }
};
typedef IceUtil::Handle<FlushCallback> FlushCallbackPtr;

void flushConnection(const Ice::ConnectionPtr& conn)
{
    FlushCallbackPtr f = new FlushCallback;
    Ice::Callback_Connection_flushBatchRequestsPtr cb =
        Ice::newCallback_Connection_flushBatchRequests(
            f, &FlushCallback::exception, &FlushCallback::sent);
    conn->begin_flushBatchRequests(cb);
}
```

For more information on asynchronous invocations, please see the relevant language mapping chapter.

The `Endpoint` Interface

The `Connection::getEndpoint` operation returns an interface of type `Endpoint`:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    const short TCPEndpointType = 1;
    const short UDPEndpointType = 3;

    local class EndpointInfo {
        int timeout;
        bool compress;
        short type();
        bool datagram();
        bool secure();
    };

    local interface Endpoint {
        EndpointInfo getInfo();
        string toString();
    };

    local class IPEndpointInfo extends EndpointInfo {
        string host;
        int port;
        string sourceAddress;
    };

    local class TCPEndpointInfo extends IPEndpointInfo {};

    local class UDPEndpointInfo extends IPEndpointInfo {
        byte protocolMajor;
        byte protocolMinor;
        byte encodingMajor;
        byte encodingMinor;
        string mcastInterface;
        int mcastTtl;
    };

    local class WSEndpointInfo extends IPEndpointInfo {
        string resource;
    };

    local class OpaqueEndpointInfo extends EndpointInfo {
        Ice::EncodingVersion rawEncoding;
        Ice::ByteSeq rawBytes;
    };
};

module IceSSL {
    const short EndpointType = 2;

    local class EndpointInfo extends Ice::IPEndpointInfo {};
};
```


Copyright 2016, ZeroC, Inc.

The `getInfo` operation returns an `EndpointInfo` instance. Note that the object returned by `getInfo` implements a more derived interface, depending on the endpoint type. You can down-cast the returned class instance and access the endpoint-specific information according to the type of the endpoint, as returned by the `type` operation.

The `timeout` member provides the timeout in milliseconds. The `compress` member is true if the endpoint uses compression (if available). The `datagram` operation returns true if the endpoint is for a datagram transport, and the `secure` operation returns true if the endpoint uses SSL.

The derived classes provide further detail about the endpoint according to its type.

Opaque Endpoints

An application may receive a proxy that contains an endpoint whose type is unrecognized by the Ice run time. In this situation, Ice preserves the endpoint in its encoded (*opaque*) form so that the proxy remains intact, but Ice ignores the endpoint for all connection-related activities. Preserving the endpoint allows an application to later forward that proxy with all of its original endpoints to a different program that might support the endpoint type in question.

Although a connection will never return an opaque endpoint, it is possible for a program to encounter an opaque endpoint when iterating over the endpoints returned by the proxy method `ice_getEndpoints`.

As a practical example, consider a program for which the IceSSL plug-in is not configured. If this program receives a proxy containing an SSL endpoint, Ice treats it as an opaque endpoint such that calling `getInfo` on the endpoint object returns an instance of `OpaqueEndpointInfo`.

Note that the `type` operation of the `OpaqueEndpointInfo` object returns the *actual* type of the endpoint. For example, the operation returns the value 2 if the object encodes an SSL endpoint. As a result, your program cannot assume that an `EndpointInfo` object whose type is 2 can be safely down-cast to `IceSSL::EndpointInfo`; if the IceSSL plug-in is not configured, such a down-cast will fail because the object is actually an instance of `OpaqueEndpointInfo`.

Client-Side Connection Usage

Clients obtain a connection by using one of the proxy methods `ice_getConnection` or `ice_getCachedConnection`. If the proxy does not yet have a connection, the `ice_getConnection` method immediately attempts to establish one. As a result, the caller must be prepared for this method to block and raise connection failure exceptions. (Use the asynchronous version of this method to avoid blocking.) If the proxy denotes a collocated object and collocation optimization is enabled, calling `ice_getConnection` returns null.

If you wish to obtain the proxy's connection without the potential for triggering connection establishment, call `ice_getCachedConnection`; this method returns null if the proxy is not currently associated with a connection or if connection caching is disabled for the proxy.

As an example, the C++ code below illustrates how to obtain a connection from a proxy and print its type:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx proxy = ...
try
{
    Ice::ConnectionPtr conn = proxy->ice_getConnection();
    if(conn)
    {
        cout << conn->type() << endl;
    }
    else
    {
        cout << "collocated" << endl;
    }
}
catch(const Ice::LocalException& ex)
{
    cout << ex << endl;
}
```

Server-Side Connection Usage

Servers can access a connection via the `con` member of the `Ice::Current` parameter passed to every operation. For collocated invocations, `con` has a nil value.

For example, this Java code shows how to invoke `toString` on the connection:

Java

```
public int add(int a, int b, Ice.Current curr)
{
    if (curr.con != null)
    {
        System.out.println("Request received on connection:\n" +
curr.con.toString());
    }
    else
    {
        System.out.println("collocated invocation");
    }
    return a + b;
}
```

Although the mapping for the Slice operation `toString` results in a Java method named `_toString`, the Ice run time implements `toStrin g` to return the same value.

Closing a Connection

Applications should rarely need to close a connection explicitly, but those that do must be aware of its implications. Since there are two ways

Copyright 2016, ZeroC, Inc.

to close a connection, we discuss them separately.

Graceful Closure

Passing an argument of `false` to the `close` operation initiates graceful connection closure, as discussed in Connection Closure. The operation blocks until all pending outgoing requests on the connection have completed.

Forceful Closure

A forceful closure is initiated by passing an argument of `true` to the `close` operation, causing the peer to receive a `ConnectionLostException`.

A client must use caution when forcefully closing a connection. Any outgoing requests that are pending on the connection when `close` is invoked will fail with a `ForcedCloseConnectionException`. Furthermore, requests that fail with this exception are not automatically retried.

In a server context, forceful closure can be useful as a defense against hostile clients.

The Ice run time interprets a `CloseConnectionException` to mean that it is safe to retry the request without violating at-most-once semantics. If automatic retries are enabled, a client must only initiate a graceful close when it knows that there are no outgoing requests in progress on that connection, or that any pending requests can be safely retried.

See Also

- The Current Object
- Automatic Retries
- Connection Establishment
- Connection Closure
- Bidirectional Connections
- IceSSL

Copyright 2016, ZeroC, Inc.

Connection Closure

The Ice run time may close a connection for many reasons, including the situations listed below:

- When deactivating an object adapter or shutting down a communicator
- As required by active connection management
- When initiated by an application
- After a network operation times out
- In response to an exception, such as a socket failure or protocol error

In most cases, the Ice run time closes a connection gracefully as required by the Ice protocol. The Ice run time only closes a connection forcefully when a timeout or protocol error occurs, or when the application explicitly requests it.

On this page:

- Graceful Connection Closure
- Connection Closure and Oneway Invocations

Graceful Connection Closure

Gracefully closing a connection occurs in stages:

- In the process that initiates closure, incoming and outgoing requests that are in progress are allowed to complete, and then a close connection message is sent to the peer. Any incoming requests received after closure is initiated are silently discarded (but may be retried, as discussed in the next bullet). An attempt to make a new outgoing request on the connection results in a `CloseConnectionException` and an automatic retry (if enabled).
- Upon receipt of a close connection message, the Ice run time in the peer closes its end of the connection. Any outgoing requests still pending on that connection fail with a `CloseConnectionException`. This exception indicates to the Ice run time that it is safe to retry those requests without violating at-most-once semantics, assuming automatic retries have not been disabled.
- After detecting that the peer has closed the connection, the initiating Ice run time closes the connection.

Connection Closure and Oneway Invocations

Oneway invocations are generally considered reliable because they are sent over a stream-oriented transport. However, it is quite possible for oneway requests to be silently discarded if a server has initiated graceful connection closure. Whereas graceful closure causes a discarded twoway request to receive a `CloseConnectionException` and eventually be retried, the sender receives no notice about a discarded oneway request.

If an application makes assumptions about the reliability of oneway requests, it may be necessary to control the events surrounding connection closure as much as possible, for example by disabling active connection management and avoiding explicit connection closures.

See Also

- Active Connection Management
- Using Connections
- Protocol Messages
- Connection Establishment
- Oneway Invocations
- Automatic Retries

Copyright 2016, ZeroC, Inc.

Bidirectional Connections

On this page:

- Use Cases for Bidirectional Connections
- Configuring a Client for Bidirectional Connections
  - Active Connection Management Considerations
- Configuring a Server for Bidirectional Connections
- Fixed Proxies
- Limitations of Bidirectional Connections
- Threading Considerations for Bidirectional Connections

Use Cases for Bidirectional Connections

An Ice connection normally allows requests to flow in only one direction. If an application's design requires the server to make callbacks to a client, the server usually establishes a new connection to that client in order to send callback requests, as shown below:

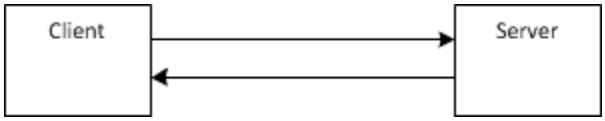

Callbacks in an open network.

Unfortunately, network restrictions often prevent a server from being able to create a separate connection to the client, such as when the client resides behind a firewall as shown here:

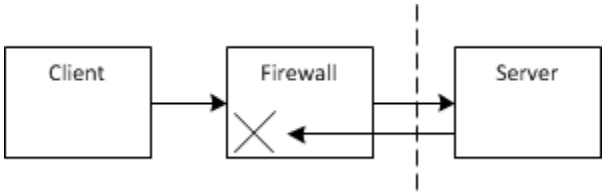

Callbacks with a firewall.

In this scenario, the firewall blocks any attempt to establish a connection directly to the client.

For situations such as these, a bidirectional connection offers a solution. Requests may flow in both directions over a bidirectional connection, enabling a server to send callback requests to a client over the client's existing connection to the server.

There are two ways to make use of a bidirectional connection. First, you can use a Glacier2 router, in which case bidirectional connections are used automatically. If you do not require the functionality offered by Glacier2 or you do not want an intermediary service between clients and servers, you can configure bidirectional connections manually.

The remainder of this section discusses manual configuration of bidirectional connections. An example that demonstrates how to configure a bidirectional connection is provided in the directory `demo/Ice/bidir` of your Ice distribution.

Configuring a Client for Bidirectional Connections

A client needs to perform the following steps in order to configure a bidirectional connection:

1. Create an object adapter to receive callback requests. This adapter does not require a name or endpoints if its only purpose is to receive callbacks over bidirectional connections.
2. Register the callback object with the object adapter.
3. Activate the object adapter.
4. Obtain a connection object by calling `ice_getConnection` on the proxy.
5. Invoke `setAdapter` on the connection, passing the callback object adapter. This associates an object adapter with the connection and enables callback requests to be dispatched.
6. Pass the identity of the callback object to the server.

The C++ code below illustrates these steps:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectAdapterPtr adapter = communicator->createObjectAdapter("");
Ice::Identity ident;
ident.name = IceUtil::generateUUID();
ident.category = "";
CallbackPtr cb = new CallbackI;
adapter->add(cb, ident);
adapter->activate();
proxy->ice_getConnection()->setAdapter(adapter);
proxy->addClient(ident);
```

The last step may seem unusual because a client would typically pass a proxy to the server, not just an identity. For example, you might be tempted to give the proxy returned by the adapter operation `add` to the server, but this would not have the desired effect: if the callback object adapter is configured with endpoints, the server would attempt to establish a separate connection to one of those endpoints, which defeats the purpose of a bidirectional connection. It is just as likely that the callback object adapter has no endpoints, in which case the proxy is of no use to the server.

Similarly, you might try invoking `createProxy` on the connection to obtain a proxy that the server can use for callbacks. This too would fail, because the proxy returned by the connection is for local use only and cannot be used by another process.

As you will see below, the server must create its own callback proxy.

Active Connection Management Considerations

Active Connection Management (ACM) automatically and transparently closes idle connections. As far as the Ice run time is concerned, an outgoing connection is a client connection regardless of whether it's also used as a bidirectional connection, therefore the client-side ACM properties govern the ACM behavior for bidirectional connections in a client. Also note that the client's outgoing connection is the only channel on which the server can send callback invocations to the client, therefore prematurely closing the connection (such as allowing the ACM facility to close it automatically) introduces the risk that the client might unknowingly fail to receive callbacks.

As of Ice 3.6, it's not necessary to disable client-side ACM when using bidirectional connections. If you leave ACM enabled, you simply need to ensure that the connection remains active to prevent the ACM facility from closing it. The easiest way to accomplish this is to enable connection heartbeats:

```
Ice.ACM.Client.Heartbeat=Always
```

Alternatively, you can enable heartbeats on a bidirectional connection by calling `setACM`:

C++

```
proxy->ice_getCachedConnection()->setACM(IceUtil::None, IceUtil::None,
Ice::HeartbeatAlways);
```

The server's ACM configuration also plays an important role here. For more information, refer to our discussion of ACM configurations for bidirectional connections.

Configuring a Server for Bidirectional Connections

A server needs to take the following steps in order to make callbacks over a bidirectional connection:

1. Obtain the identity of the callback object, which is typically supplied by the client.
2. Create a proxy for the callback object by calling `createProxy` on the connection. The connection object is accessible as a member of the `Ice::Current` parameter to an operation implementation.

Copyright 2016, ZeroC, Inc.

These steps are illustrated in the C++ code below:

C++

```
void addClient(const Ice::Identity& ident, const Ice::Current& curr)
{
    CallbackPrx client = CallbackPrx::uncheckedCast(curr.con->createProxy(ident));
    client->notify();
}
```

Fixed Proxies

The proxy returned by a connection's `createProxy` operation is called a *fixed proxy*. It can only be used in the server process and cannot be marshaled or stringified by `proxyToString`; attempts to do so raise `FixedProxyException`.

A fixed proxy is bound to the connection that created it, and ceases to work once that connection is closed. If the connection is closed prematurely, either by active connection management (ACM) or by explicit action on the part of the application, the server can no longer make callback requests using that proxy. Any attempt to use the proxy again usually results in a `CloseConnectionException`.

Many aspects of a fixed proxy cannot be changed. For example, it is not possible to change the proxy's endpoints or timeout. Attempting to invoke a method such as `ice_timeout` on a fixed proxy raises `FixedProxyException`.

Limitations of Bidirectional Connections

Bidirectional connections have certain limitations:

- They can only be configured for connection-oriented transports such as TCP and SSL.
- Most proxy factory methods are not relevant for a proxy created by a connection's `createProxy` operation. The proxy is bound to an existing connection, therefore the proxy reflects the connection's configuration. Attempting to change settings such as the proxy's timeout value causes the Ice run time to raise `FixedProxyException`. Note however that it is legal to configure a fixed proxy for using oneway or twoway invocations. You may also invoke `ice_secure` on a fixed proxy if its security configuration is important; a fixed proxy configured for secure communication raises `NoEndpointException` on the first invocation if the connection is not secure.
- A connection established from a Glacier2 router to a server is not configured for bidirectional use. Only the connection from a client to the router is bidirectional. However, the client must not attempt to manually configure a bidirectional connection to a router, as this is handled internally by the Ice run time.

Threading Considerations for Bidirectional Connections

The Ice run time normally creates two thread pools for processing network traffic on connections: the client thread pool manages outgoing connections and the server thread pool manages incoming connections. All of the object adapters in a server share the same thread pool by default, but an object adapter can also be configured to have its own thread pool. The default size of the client and server thread pools is one.

The client thread pool processes replies to pending requests. When a client configures an outgoing connection for bidirectional requests, the client thread pool also becomes responsible for dispatching callback requests received over that connection.

Similarly, the server thread pool normally dispatches requests from clients. If a server uses a bidirectional connection to send callback requests, then the server thread pool must also process the replies to those requests.

You must increase the size of the appropriate thread pool if you need the ability to dispatch multiple requests in parallel, or if you need to make nested twoway invocations. For example, a client that receives a callback request over a bidirectional connection and makes nested invocations must increase the size of the *client* thread pool.

See Also

- Glacier2
- Creating an Object Adapter
- Servant Activation and Deactivation
- Using Connections
- Object Identity

Copyright 2016, ZeroC, Inc.

- Active Connection Management
- Proxy Methods
- Nested Invocations
- The Ice Threading Model
- Object Adapter Thread Pools

Copyright 2016, ZeroC, Inc.

Connection Timeouts

On this page:

- Overview of Connection Timeouts
- Configuring Connection Timeouts
- Connection Timeout Failures
- ACM and Timeouts
- Connection Reuse and Timeouts

Overview of Connection Timeouts

As of Ice 3.6, connection timeouts only affect network operations. Use invocation timeouts to limit the amount of time a client waits for an operation to complete.

Connection timeouts allow applications to detect low-level network problems in a reasonable period of time. Ice enforces connection timeouts when performing network operations such as establishing a connection, reading and writing to a connection, and closing a connection. Disabling connection timeouts means an application may not discover a network issue until much later (if at all) when low-level network protocols finally detect and report the problem, therefore Ice enables connection timeouts by default and we strongly encourage you to use them.

You should normally choose your connection timeouts based on the speed of the network on which the connections take place. For example, the connection timeout for a local gigabit network will usually be much smaller than the timeout for a 56kbps connection.

Finally, note that timeouts in Ice are "soft" timeouts, in the sense that they are not precise, real-time timeouts. (The precision is limited by the capabilities of the underlying operating system.)

Configuring Connection Timeouts

Ice supports several configuration properties that you can use to control connection timeouts:

- `Ice.Default.Timeout`
  Added in Ice 3.6, this property specifies the default timeout in milliseconds for all connections. This property's default value of `60000` means connection timeouts are enabled by default with a timeout of 60 seconds.
- `Ice.Override.ConnectTimeout`
  This setting overrides any existing connection timeout, but only applies while Ice establishes a connection.
- `Ice.Override.CloseTimeout`
  This setting overrides any existing connection timeout, but only applies while Ice closes a connection.
- `Ice.Override.Timeout`
  This setting overrides any existing connection timeout. It applies to all network operations unless superseded by `Ice.Override.ConnectTimeout` or `Ice.Override.CloseTimeout`.

You can also configure connection timeouts individually for the endpoints of a proxy. Consider this example:

```
MyProxy=hello:tcp -h 10.0.0.1 -t 1000:tcp -h 205.125.53.4 -t 5000
```

This stringified proxy contains two TCP endpoints. The first endpoint refers to a host in a local network and its `-t` option specifies a connection timeout of one second. The second endpoint refers to a host on the internet and uses a connection timeout of five seconds. If a proxy's endpoints don't include timeouts, the endpoints inherit the communicator's default timeout set by `Ice.Default.Timeout`. The various `Ice.Override` properties take precedence over any endpoint timeouts.

Finally, the proxy method `ice_timeout` returns a new proxy in which all of its endpoints have the specified timeout:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx p = communicator->stringToProxy("hello:tcp -h 10.0.0.1 -t 1000:tcp
-h 205.125.53.4 -t 5000");

cout << p->ice_toString() << endl; // Output: hello:tcp -h 10.0.0.1 -t 1000:tcp -h
205.125.53.4 -t 5000
p = p->ice_timeout(1500);
cout << p->ice_toString() << endl; // Output: hello:tcp -h 10.0.0.1 -t 1500:tcp -h
205.125.53.4 -t 1500
```

The `Ice.Override` properties also take precedence over any timeouts configured via `ice_timeout`.

Connection Timeout Failures

Ice considers a connection timeout to indicate a serious error with the underlying network connection, causing Ice to immediately close the connection and report a `TimeoutException`, or one of its subclasses `ConnectTimeoutException` or `CloseTimeoutException` if appropriate, for all of the invocations pending on that connection.

The application may not receive an exception immediately after a connection timeout occurs, as Ice may attempt the invocation again if automatic retries are enabled and the invocation qualifies for retry. Since its default configuration allows Ice to perform one retry without a delay, it will appear to the application as if the timeout takes twice as long as expected to be reported.

In most cases you won't need to handle a timeout as a special case and can use the typical exception clauses:

C++

```
Filesystem::FilePrx myFile = ...;
myFile = myFile->ice_timeout(2500);

try {
    Lines text = myFile->read();
} catch (const Filesystem::GenericError & ex) {
    cerr << ex.reason << endl;
} catch (const Ice::LocalException & ex) {
    cerr << "failed: " << ex << endl;
}
```

If your application needs to treat timeouts differently, you can catch them as follows:

Copyright 2016, ZeroC, Inc.

C++

```
Filesystem::FilePrx myFile = ...;
myFile = myFile->ice_timeout(2500);

try {
    Lines text = myFile->read();
} catch (const Filesystem::GenericError & ex) {
    cerr << ex.reason << endl;
} catch (const Ice::ConnectTimeoutException &) {
    cerr << "connect timed out!" << endl;
} catch (const Ice::CloseTimeoutException &) {
    cerr << "close timed out!" << endl;
} catch (const Ice::TimeoutException &) {
    cerr << "timed out!" << endl;
} catch (const Ice::LocalException & ex) {
    cerr << "failed: " << ex << endl;
}
```

ACM and Timeouts

The new Active Connection Management features in Ice 3.6 can also help your application in detecting and dealing with connectivity issues and complement the connection timeout behavior. For example, connection timeouts only have an effect when the Ice run time is actively performing network operations, but it's also possible for problems to arise during periods of inactivity. These problems may not be noticed for some time.

The ACM heartbeat facility can optionally be configured to send heartbeat messages at regular intervals when a connection would otherwise be idle. Doing so makes it more likely that the Ice run time will detect a connectivity issue in a timely manner.

Connection Reuse and Timeouts

Ice tries to reuse existing connections as much as possible unless the application indicates otherwise. When an application makes its first invocation on a proxy, Ice searches its list of open connections to see if any of them match any of the proxy's endpoints. Determining whether an existing connection is "compatible" with a proxy endpoint, and therefore suitable for reuse, involves several criteria. One attribute that Ice considers is the timeout: the endpoint's timeout and the connection's timeout must match for the connection to be considered eligible for reuse.

If you encounter a situation where Ice is opening a new connection when you expect it to use an existing connection, your timeout configuration may be the reason.

See Connection Establishment for more information on how Ice decides whether to open a new connection or reuse an existing one.

See Also

- Proxy Methods
- Invocation Timeouts
- Automatic Retries
- Connection Establishment

Copyright 2016, ZeroC, Inc.

Collocated Invocation and Dispatch

On this page:

- Overview
- Collocated Invocations
- Creating an Object Adapter for Collocated Invocations

Overview

One of the useful features of the Ice run time is that it is *location transparent*: the client does not need to know where the implementation of an Ice object resides; an invocation on an object automatically is directed to the correct target, whether the object is implemented in the local address space, in another address space on the same machine, or in another address space on a remote machine. Location transparency is important because it allows us to change the location of an object implementation without breaking client programs and, by using a location service such as IceDiscovery or IceGrid, addressing information such as host names and port numbers can be externalized so they do not appear in stringified proxies.

Collocated Invocations

For invocations that cross address space boundaries (or more accurately, cross communicator boundaries), the Ice run time dispatches requests via the appropriate transport. However, for a proxy invocation in which the proxy and the servant that processes the invocation share the same communicator (so-called *collocated* invocations), the Ice run time, by default, does not send the invocation via the transport specified in the proxy. Instead, collocated invocations take a short-cut inside the Ice run time and are dispatched more efficiently.

Note that if the proxy and the servant do not use the same communicator, the invocation is *not* collocated, even though caller and callee are in the same address space.

The reason for this is if collocated invocations were sent via TCP/IP, for example, invocations would still be sent via the operating system kernel (using the back plane instead of a network) and would incur the full cost of creating TCP/IP connections, trapping in and out of the kernel, and so on. By optimizing collocated requests, much of this overhead can be avoided.

For efficiency reasons, collocated invocations are not completely location transparent, that is, a collocated call has semantics that differ in some ways from calls that cross address-space boundaries. Specifically, collocated invocations differ from ordinary invocations in the following respects:

- Most collocated invocations are dispatched in the server-side thread pool just like regular invocations; the only exceptions are synchronous twoway collocated invocations with no invocation timeout, which are dispatched in the calling thread.
- The object adapter holding state is ignored: collocated invocations proceed normally even if the target object's adapter is in the holding state.
- AMI callbacks for asynchronous collocated invocations are dispatched from the servant's calling thread and not from the client-side thread pool unless AMD is used for the servant dispatch. In this case, the AMI callback is called from the client-side thread pool.
- Invocation timeouts work as usual, but connection timeouts are ignored.

In practice, these differences rarely matter. The most likely cause of surprises with collocated invocations is dispatch in the calling thread, that is, a collocated invocation behaves like a local, synchronous procedure call. This can cause problems if, for example, the calling thread acquires a lock that an operation implementation tries to acquire as well: unless you use recursive mutexes, this will cause deadlock.

The Ice run time uses the following semantics to determine whether a proxy is eligible for the collocated optimization:

- For an indirect proxy, collocation optimization is used if the proxy's adapter ID matches the adapter ID or replica group ID of an object adapter in the same communicator.
- For a well-known proxy, the Ice run time queries each object adapter to determine if the servant is local.
- For a direct proxy, the Ice run time performs an endpoint search using the proxy's endpoints.

When an endpoint search is required, the Ice run time compares each of the proxy's endpoints against the endpoints of the communicator's object adapters. Only the transport, address and port are considered; other attributes of an endpoint, such as timeout settings, are not considered during this search. If a match is found, the invocation is dispatched using collocation optimization. Normally this search is executed only once, during the proxy's first invocation, although the proxy's connection caching setting influences this behavior.

Collocation optimization is enabled by default, but you can disable it for all proxies by setting the property `Ice.Default.CollocationOptimized=0`. You can also disable the optimization for an individual proxy using the factory method `ice_collocationOptimized(false)`. Finally, for proxies created from a property using `propertyToProxy`, the property *`name`*`.CollocationOptimized` configures the default setting for the proxy.

Copyright 2016, ZeroC, Inc.

Creating an Object Adapter for Collocated Invocations

An object adapter requires no endpoints if its only purpose is to dispatch collocated requests. You can create a "collocated-only" object adapter by calling `createObjectAdapter` with an empty name:

C++

```
Ice::ObjectAdapterPtr collocAdapter = communicator->createObjectAdapter("");
```

Normally Ice requires an object adapter to either have endpoints or be configured with a router, but Ice relaxes this restriction when the name is empty. Note however that this means collocated-only object adapters cannot be configured using properties.

Proxies created by a collocated-only object adapter contain no endpoints, which essentially makes them well-known proxies. When the program makes its initial invocation on such a proxy, Ice searches every local object adapter for a servant with an identity matching that of the proxy. For Ice to successfully find the collocated servant, it must be present in the object adapter's active servant map. Servant locators and default servants are not queried in this situation.

See Also

- IceDiscovery
- IceGrid
- Threads and Concurrency with C++
- Connection Establishment
- Proxy Methods
- Obtaining Proxies

Copyright 2016, ZeroC, Inc.

Locators

In Terminology, we described briefly how the Ice run time uses an intermediary, known as a *location service*, to convert the symbolic information in an indirect proxy into an endpoint that it can use to communicate with a server. This section expands on that introduction to explain in more detail how the Ice run time interacts with a location service. You can create your own location service or you can use IceDiscovery or IceGrid, which are both implementations of a location service. Describing how to implement a location service is outside the scope of this manual.

A *locator* is an Ice object that is implemented by a location service. A locator object must support the Slice interface `Ice::Locator`, which defines operations that satisfy the location requirements of the Ice run time. Applications do not normally use these operations directly, but the locator object may support an implementation-specific interface derived from `Ice::Locator` that provides additional functionality. For example, IceGrid's locator object provides access to an `IceGrid::Query` object so that applications can perform more sophisticated queries.

Topics

- Locator Semantics for Clients
- Locator Configuration for a Client
- Locator Semantics for Servers
- Locator Configuration for a Server

See Also

- Terminology
- IceDiscovery
- IceGrid
- Querying Well-Known Objects

Copyright 2016, ZeroC, Inc.

Locator Semantics for Clients

On this page:

- Invocations with an Indirect Proxy
- Replication with a Locator
- Locator Cache
- Locator Cache Timeout
- Load Balancing with a Locator

Invocations with an Indirect Proxy

On the first use of an indirect proxy in an application, the Ice run time may issue a remote invocation on the locator object. This activity is transparent to the application, as shown below:

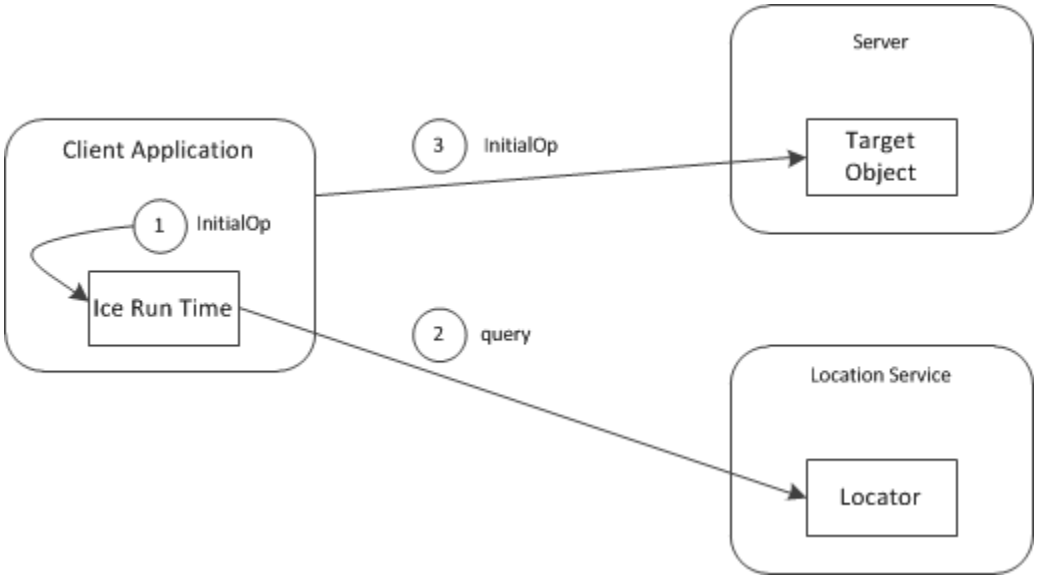

Locating an object.

1. The client invokes the operation `initialOp` on an indirect proxy.
2. The Ice run time checks an internal cache (called the *locator cache*) to determine whether a query has already been issued for the symbolic information in the proxy. If so, the cached endpoint is used and an invocation on the locator object is avoided. Otherwise, the Ice run time issues a locate request to the locator.
3. If the object is successfully located, the locator returns its current endpoints. The Ice run time in the client caches this information, es tablishes a connection to one of the endpoints, and proceeds to send the invocation as usual.
4. If the object's endpoints cannot be determined, the client receives an exception. `NotRegisteredException` is raised when an identity, object adapter identifier or replica group identifier is not known. A client may also receive `NoEndpointException` if the location service failed to determine the current endpoints.

As far as the Ice run time is concerned, the locator simply converts the information in an indirect proxy into usable endpoints. Whether the locator's implementation is more sophisticated than a simple lookup table is irrelevant to the Ice run time. However, the act of performing this conversion may have additional semantics that the application must be prepared to accept.

For example, when using IceGrid as your location service, the target server may be launched automatically if it is not currently running, and the locate request does not complete until that server is started and ready to receive requests. As a result, the initial request on an indirect proxy may incur additional overhead as all of this activity occurs.

Replication with a Locator

An indirect proxy may substitute a replica group identifier in place of the object adapter identifier. In fact, the Ice run time does not distinguish between these two cases and considers a replica group identifier as equivalent to an object adapter identifier for the purposes of resolving the proxy. The location service implementation must be able to distinguish between replica groups and object adapters using only this identifier.

The location service may return multiple endpoints in response to a locate request for an adapter or replica group identifier. These endpoints might all correspond to a single object adapter that is available at several addresses, or to multiple object adapters each listening at a single

Copyright 2016, ZeroC, Inc.

address, or some combination thereof. The Ice run time attaches no semantics to the collection of endpoints, but the application can make assumptions based on its knowledge of the location service's behavior.

When a location service returns more than one endpoint, the Ice run time behaves exactly as if the proxy had contained several endpoints. As always, the goal of the Ice run time is to establish a connection to one of the endpoints and deliver the client's request. By default, all requests made via the proxy that initiated the connection are sent to the same server until that connection is closed.

After the connection is closed, such as by Active Connection Management (ACM), subsequent use of the proxy causes the Ice run time to obtain another connection. Whether that connection uses a different endpoint than previous connections depends on a number of factors, but it is possible for the client to connect to a different server than for previous requests.

Locator Cache

After successfully resolving an indirect proxy, the location service must return at least one endpoint. How the service derives the list of endpoints that corresponds to the proxy is entirely implementation dependent. For example, IceGrid's location service can be configured to respond in a variety of ways; one possibility uses a simple round-robin scheme, while another selects endpoints based on the system load of the target hosts.

A locate request has the potential to significantly increase the latency of the application's invocation with a proxy, and this is especially true if the locate request triggers additional implicit actions such as starting a new server process. Fortunately, this overhead is normally incurred only during the application's initial invocation on the proxy, but this impact is influenced by the Ice run time's caching behavior.

To minimize the number of locate requests, the Ice run time caches the results of previous requests. By default, the results are cached indefinitely, so that once the Ice run time has obtained the endpoints associated with an indirect proxy, it never issues another locate request for that proxy. Furthermore, the default behavior of a proxy is to cache its connection, that is, once a proxy has obtained a connection, it continues to use that connection indefinitely.

Taken together, these two caching characteristics represent the Ice run time's best efforts to optimize an application's use of a location service: after a proxy is associated with a connection, all future invocations on that proxy are sent on the same connection without any need for cache lookups, locate requests, or new connections.

If a proxy's connection is closed, the next invocation on the proxy prompts the Ice run time to consult its locator cache to obtain the endpoints from the prior locate request. Next, the Ice run time searches for an existing connection to any of those endpoints and uses that if possible, otherwise it attempts to establish a new connection to each of the endpoints until one succeeds. Only if that process fails does the Ice run time clear the entry from its cache and issue a new locate request with the expectation that a usable endpoint is returned.

The Ice run time's default behavior is optimized for applications that require minimal interaction with the location service, but some applications can benefit from more frequent locate requests. Normally this is desirable when implementing a load-balancing strategy, as we discuss in more detail below. In order to increase the frequency of locate requests, an application must configure a timeout for the locator cache and manipulate the connections of its proxies.

Locator Cache Timeout

An application can define a timeout to control the lifetime of entries in the locator cache. This timeout can be specified globally using the `Ice.Default.LocatorCacheTimeout` property and for individual proxies using the proxy method `ice_locatorCacheTimeout`. The Ice run time's default behavior is equivalent to a timeout value of `-1`, meaning the cache entries never expire. Using a timeout value greater than zero causes the cache entries to expire after the specified number of seconds. Finally, a timeout value of zero disables the locator cache altogether.

The previous section explained the circumstances in which the Ice run time consults the locator cache. Briefly, this occurs only when the application has invoked an operation on a proxy and the proxy is not currently associated with a connection. If the timeout is set to zero, the Ice run time issues a new locate request immediately. Otherwise, for a non-zero timeout, the Ice run time examines the locator cache to determine whether the endpoints from the previous locate request have expired. If so, the Ice run time discards them and issues a new locate request.

Given this behavior, if your goal is to force a proxy invocation to issue locate requests more frequently, you can do so only when the proxy is not associated with a connection. You can accomplish that in several ways:

- create a new proxy, which is inherently not connected by default
- explicitly close the proxy's existing connection
- disabling the proxy's connection caching behavior

Of these choices, the last is the most common.

Load Balancing with a Locator

Copyright 2016, ZeroC, Inc.

Ice supports proxy-based load balancing whose behavior is driven solely by a proxy's configuration settings. A disadvantage of relying solely on this form of load balancing is that the client cannot make any intelligent decisions based on the status of the servers. If you want to distribute your requests in a more sophisticated way, you must either modify your clients to query the servers directly, or use a location service that can transparently direct a client to an appropriate server. For example, the IceGrid location service can monitor the system load on each server host and use that information when responding to locate requests.

The location service may return only one endpoint, which presumably represents the best server (at that moment) for the client to use. With only one endpoint available, changing the proxy's endpoint selection type makes no difference. However, by disabling connection caching and modifying the locator cache timeout, the application can force the Ice run time to periodically retrieve an updated endpoint from the location service. For example, an application can set a locator cache timeout of thirty seconds and communicate with the selected server for that period. After the timeout has expired, the next invocation prompts the Ice run time to issue a new locate request, at which point the client might be directed to a different server.

If the location service returns multiple endpoints, the application must be designed with knowledge of how to interpret them. For instance, the location service may attach semantics to the order of the endpoints (such as least-loaded to most-loaded) and intend that the application use the endpoints in the order provided. Alternatively, the client may be free to select any of the endpoints. As a result, the application and the location service must cooperate to achieve the desired results.

You can combine the simple form of load balancing described in the previous section with an intelligent location service to gain even more flexibility. For example, suppose an application expects to receive multiple endpoints from the location service and has configured its proxy to disable connection caching and set a locator cache timeout. For each invocation, Ice run time selects one of the endpoints provided by the location service. When the timeout expires, the Ice run time issues a new locate request and obtains a fresh set of endpoints from which to choose.

See Also

- Terminology
- Proxy Methods
- Connection Establishment
- Active Connection Management
- Ice.Default.*
- Proxy-Based Load Balancing

Copyright 2016, ZeroC, Inc.

Locator Configuration for a Client

An Ice client application must supply a proxy for the locator object, which it can do in several ways:

- by explicitly configuring an indirect proxy using the `ice_locator` proxy method
- by calling `setDefaultLocator` on a communicator, after which all new proxies use the given locator by default
- by defining the `Ice.Default.Locator` configuration property, which causes all proxies to use the given locator by default

The Ice run time's efforts to resolve an indirect proxy can be traced by setting the following configuration properties:

```
Ice.Trace.Network=2
Ice.Trace.Protocol=1
Ice.Trace.Locator=2
```

See Ice.Trace.* for more information on these properties.

See Also

- Locator Semantics for Clients
- Proxy Methods
- Ice.Default.*
- Ice.Trace.*

Copyright 2016, ZeroC, Inc.

Locator Semantics for Servers

A location service must know the endpoints of any object adapter whose identifier can be used in an indirect proxy. For example, suppose a client uses the following proxy:

```
Object1@PublicAdapter
```

The Ice run time in the client includes the identifier `PublicAdapter` in its locate request and expects to receive the associated endpoints. The only way the location service can know these endpoints is if it is given them. When you consider that an object adapter's endpoints may not specify fixed ports, and therefore the endpoint addresses may change each time the object adapter is activated, it is clear that the best source of endpoint information is the object adapter itself. As a result, an object adapter that is properly configured contacts the locator during activation to supply its identifier and current endpoints. More specifically, the object adapter registers itself with an object implementing the `Ice::LocatorRegistry` interface, whose proxy the object adapter obtains from the locator.

A location service may require that all object adapters be pre-registered via some implementation-specific mechanism. (IceGrid behaves this way by default.) This implies that activation can fail if the object adapter supplies an identifier that is unknown to the location service. In such a situation, the object adapter's `activate` operation raises `NotRegisteredException`.

In a similar manner, an object adapter that participates in a replica group includes the group's identifier in the locator request that is sent during activation. If the location service requires replica group members to be configured in advance, `activate` raises `NotRegisteredException` if the object adapter's identifier is not one of the group's registered participants.

See Also

- Object Adapters
- Locator Semantics for Clients
- Locator Configuration for a Server
- IceGrid
- Terminology
- Object Adapter Replication

Copyright 2016, ZeroC, Inc.

Locator Configuration for a Server

On this page:

- Configuring an Object Adapter with a Locator
- Registering a Process with a Locator

Configuring an Object Adapter with a Locator

An object adapter must be able to obtain a locator proxy in order to register itself with a location service. Each object adapter can be configured with its own locator proxy by defining its `Locator` property, as shown in the example below for the object adapter named `SampleAdapter`:

```
SampleAdapter.Locator=IceGrid/Locator:tcp -h locatorhost -p 10000
```

Alternatively, a server may call `setLocator` on the object adapter prior to activation. If the object adapter is not explicitly configured with a locator proxy, it uses the default locator as provided by its communicator.

Two other configuration properties influence an object adapter's interactions with a location service during activation:

- `AdapterId`
  Configuring a non-empty identifier for the `AdapterId` property causes the object adapter to register itself with the location service. A locator proxy must also be configured.
- `ReplicaGroupId`
  Configuring a non-empty identifier for the `ReplicaGroupId` property indicates that the object adapter is a member of a replica group. For this property to have an effect, `AdapterId` must also be configured with a non-empty value.

We can use these properties as shown below:

```
SampleAdapter.AdapterId=SampleAdapterId
SampleAdapter.ReplicaGroupId=SampleGroupId
SampleAdapter.Locator=IceGrid/Locator:tcp -h locatorhost -p 10000
```

Note that a location service may enforce pre-registration requirements.

Registering a Process with a Locator

An activation service, such as an IceGrid node, needs a reliable way to gracefully deactivate a server. One approach is to use a platform-specific mechanism, such as POSIX signals. This works well on POSIX platforms when the server is prepared to intercept signals and react appropriately. On Windows platforms, it works less reliably for C++ servers, and not at all for Java servers. For these reasons, Ice provides an alternative that is both portable and reliable:

Slice

```
module Ice {
interface Process {
    void shutdown();
    void writeMessage(string message, int fd);
};
};
```

The Slice interface `Process` allows an activation service to request a graceful shutdown of the server. When `shutdown` is invoked, the object implementing this interface is expected to initiate the termination of its server process. The activation service may expect the server to terminate within a certain period of time, after which it may terminate the server abruptly.

Copyright 2016, ZeroC, Inc.

One of the benefits of the Ice administrative facility is that it creates an implementation of `Process` and makes it available via an administrative object adapter, or your own object adapter. Furthermore, IceGrid automatically enables this facility on the servers that it activates.

See Also

- Object Adapters
- Locators
- Locator Configuration for a Client
- Locator Semantics for Servers
- Portable Signal Handling in C++
- The Process Facet
- Administrative Facility
- IceGrid

Copyright 2016, ZeroC, Inc.

Slicing Values and Exceptions

On this page:

- Distinguishing Objects and Servants
- Design Considerations for Exceptions
- Design Considerations for Objects
- Composition using Slices
- The Factory Requirement
- Slice Formats
  - Compact Format
  - Sliced Format
  - Choosing a Format
- Optional Objects
- Preserving Slices
- Unknown Sliced Objects

Distinguishing Objects and Servants

The term Ice object refers to a conceptual entity that is incarnated by a programming language object called a servant. An Ice object typically implements a Slice interface, but could implement a Slice class with operations instead.

In addition, applications can transfer instances of Slice classes *by value*, using classes as just another type in the data model, along with other Slice types such as structures and sequences. We refer to this usage as transferring *objects by value*, to distinguish it from the formal concept of *Ice objects*.

As a refresher on Slice syntax, consider the following definitions:

Slice

```
class ClassWithOperations
{
    int cost;

    void updateCost();
};

interface Remote
{
    ClassWithOperations getByValue();
    ClassWithOperations* getByProxy();
};
```

As its name implies, the operation `getByValue` returns a `ClassWithOperations` object by value, whereas the `getByProxy` operation returns a proxy to a remote Ice object that implements `ClassWithOperations`. Although their syntax is quite similar, the semantics of these two operations are very different. Calling `getByValue`, a client receives a *local* instance of `ClassWithOperations` on which it can access the `cost` member and call the `updateCost` method. With `getByProxy`, a client can use the returned proxy to call `updateCost` as a remote invocation but cannot access the `cost` member, which is local state for the remote servant.

Design Considerations for Exceptions

Using exceptions to report error conditions is a generally accepted idiom, given its wide support in programming languages and its convenience in simplifying a developer's error-handling responsibilities. The alternatives, such as using numeric error codes, can be unwieldy in practice and lack the flexibility that a well-designed exception hierarchy provides. Slice does not impose any particular error-handling style on developers, but if you choose to use Slice exceptions, you should understand their advantages and limitations.

Slice exception types support single inheritance, which allows you to compose hierarchies of related error conditions. Furthermore, exceptions may have any number of data members, meaning your errors can convey as little or as much detail as you like. One limitation of exceptions is that they are treated differently than other Slice data types: exceptions cannot be passed as parameters, and cannot be used as data members. The sole purpose of an exception type is to appear in the `throws` clause of an operation.

Copyright 2016, ZeroC, Inc.

As an Ice developer, you are not expected to know the intricacies of how Ice transfers exceptions from server to client. However, it can be very useful for you to understand that exceptions are composed of slices, that multiple encoding formats are available, and that passing exceptions through intermediaries may require additional effort.

Design Considerations for Objects

The motivation for using objects by value is not as simple as for exceptions. If you wish to use the exception idiom, you have no choice but to use the Slice exception model. However, an application is not forced to use objects by value. Other user-defined Slice types, such as sequences, dictionaries, and structures, may satisfy the data model requirements of many applications. Classes and structures may seem similar at first glance because they can both be used as data containers, but there are also some significant differences between them. Be sure you understand the architectural implications of classes before using them in your application. Granted, if your data model requires the use of inheritance, polymorphism, or self-referential capabilities, you really have no other choice but to use classes.

One aspect of using objects by value that must not be overlooked is overhead, both in terms of the processing effort required to encode and decode them, as well as the space consumed in messages that use them. In Ice 3.4 and earlier, objects by value were relatively expensive, likely causing many developers to use structures when classes might have been the more convenient choice. In Ice 3.5 and later, the new compact format makes objects by value much more competitive with structures. Ice 3.5 and later also make it possible to forward objects of unknown types for situations where flexibility takes precedence over compactness.

Composition using Slices

The data encoding rules for classes and exceptions rely on the concept of *slices* in which each slice corresponds to a level in the type hierarchy and contains the encoded data members defined at that level. The low-level details of the encoding rules are not relevant here, but having a basic understanding of how Ice marshals and unmarshals classes and exceptions can help you with your design decisions. Consider this example:

Slice

```
class A
{
    int i;
};

class B extends A
{
    float f;
};

class C extends B
{
    string s;
};
```

The encoded form of an instance of class `C` is summarized below:

Slice	Contents
`C`	`string s`
`B`	`float f`
`A`	`int i`

Notice that the instance is encoded in order from most-derived to least-derived. For the sake of simplicity, we have omitted many of the encoding details. For instance, we have not shown that the initial (most-derived) slice includes a type ID string such as `"::MyModule::C"` that identifies the object's type.

Copyright 2016, ZeroC, Inc.

Although we used a class hierarchy in the example above, the behavior for exceptions is very similar.

Let's add the following interface to our discussion:

Slice

```
interface I
{
    B getObject();
};
```

Now suppose a client invokes `getObject`. Let's also suppose that the servant actually returns an instance of class `C`, which is perfectly legal given that `C` derives from `B`. The Ice run time in the client knows that the formal return type of the `getObject` operation is `B` and unmarshals the object as follows:

1. Extract the type ID of the initial slice (`"::MyModule::C"`).
2. If this type is known to the client, instantiate the object, decode the state for each of its slices, and return the object.
3. If type `C` is *not* known to the client, which can occur when client and server are using different versions of the Slice definitions, the Ice run time discards the encoded data for slice `C` (also known as *slicing* the object), reads the type ID of the next slice, and repeats step 2.
4. If type `B` is also not known to the client then we have a problem. First, from a logical standpoint, the client *must* know type `B` because it is the statically-declared return type of the operation that the client just invoked. Second, we cannot slice this object any further because it would no longer be compatible with the formal signature of the operation; the returned object must at least be an instance of `B`, so we could not return an instance of `A`. In either case, the Ice run time would raise an exception.

Generally speaking, to unmarshal an instance of a class or exception, the Ice run time discards the slices of unknown types until it finds a type that it recognizes, exhausts all slices, or can no longer satisfy the formal type signature of the operation. This slicing feature allows the receiver, whose Slice definitions may be limited or outdated, to continue to function properly even when it does not recognize the most-derived type.

With this knowledge, we can now discuss new features introduced in Ice 3.5 that affect the run time's encoding rules for classes and exceptions. Namely, a choice of encoding formats and the ability to preserve the slices of unknown types. First, however, you should understand the need for object factories.

The Factory Requirement

For Slice classes that define data members but no operations, receiving an object by value is straightforward: the Ice run time examines the type ID encoded with the object, translates that (in a language-specific manner) into the corresponding generated class, and creates an instance of that class to represent the object. This happens without any help from the application. However, things get more complicated when a class defines operations. By definition, such a class is abstract, therefore Ice cannot allocate an instance of the generated abstract class. Only the application knows which derived class implements an abstract Slice class, so Ice requires the application to supply factories for translating type IDs into their programming language equivalent objects.

Ice does not require or support user-supplied factories for exceptions because exceptions cannot define operations.

The encoding for classes does not inform the receiver whether an object is abstract. As Ice processes each slice of an object, it first attempts to locate the corresponding generated class. If Ice cannot find the class, or finds the class but determines that it is abstract, the run time looks for a factory registered by the application for that type ID. If no factory is found to create the object, Ice may discard the slice as described above.

Here is some sample code to demonstrate the implementation and registration of a factory:

Copyright 2016, ZeroC, Inc.

Java

```
private static class CWOFactory implements Ice.ObjectFactory
{
    public Ice.Object create(String typeId)
    {
        assert(typeId.equals(ClassWithOperations.ice_staticId()));
        return new ClassWithOperationsImpl();
    }

    public void destroy()
    {
    }
}

private static class ClassWithOperationsImpl extends ClassWithOperations
{
    public void updateCost()
    {
        // ...
    }
}

Ice.Communicator communicator = ...;
communicator.addObjectFactory(new CWOFactory(),
ClassWithOperations.ice_staticId());
```

Our factory is only intended to create instances of the `ClassWithOperations` type we defined earlier, which is the reason for the assertion. In your own application, you can partition the factory responsibilities however you wish. For example, you can define a separate factory class for each Slice class, or use a single factory class that knows how to create multiple types. Whatever you decide, you must register a factory for each abstract class that the program has the potential to receive, and the registrations must occur before any remote invocations are made that might return one of these types.

You can also register a factory for a concrete Slice class if you wish to substitute a derived type, or simply because you want more control over the allocation of that type.

Slice Formats

Version 1.0 of the Ice encoding, which is used by all versions of Ice through 3.4, supports a single format for encoding the slices of classes and exceptions. The primary design goal for this format was the ability to support slicing at the expense of a somewhat larger encoding. The fact that the encoding embeds the type IDs of each type in an instance's type hierarchy means the size of the encoded data is affected by the symbol names and nesting depth of the Slice definitions. Consider this example:

Copyright 2016, ZeroC, Inc.

Slice

```
module TopLevelModule
{
    module IntermediateModule
    {
        class Account { ... };
        class SubcontractorAccount extends Account { ... };
    };
};
```

The encoding for an instance of `SubcontractorAccount` includes the following type ID strings:

- `::TopLevelModule::IntermediateModule::SubcontractorAccount`
- `::TopLevelModule::IntermediateModule::Account`
- `::Ice::Object`

The last ID serves as a terminator to indicate the last slice and reflects the fact that all object types ultimately derive from the base `Object` type.

As you can see, there is a non-trivial amount of overhead for passing objects by value. Note however that the encoding includes a lookup table for type IDs to minimize this overhead as much as possible, meaning each unique type ID is only encoded once per message. While sending a single instance of `SubcontractorAccount` might be relatively expensive, sending a collection of `SubcontractorAccount` objects amortizes this cost over a number of instances. We can contrive situations where a message contains many instances of unrelated types, in which case the encoded type IDs would have a greater impact on the size of the resulting message, but those scenarios are less common.

The slice format for classes and exceptions is nearly identical. The main difference is that, for an exception, there is no need for a type ID lookup table because there can never be more than one instance of an exception in a message.

With version 1.1 of the encoding introduced in Ice 3.5, you now have two slice formats to choose from: *compact* and *sliced*.

Compact Format

The compact format in version 1.1 of the encoding sacrifices the slicing feature of classes and exceptions to reduce the encoded size of these types. The format still encodes the type ID of the most-derived type, and continues to use a lookup table to prevent duplication of type IDs in a message, but omits the type IDs for any intermediate types. Version 1.1 of the encoding also eliminates the need to encode the type ID for the base `Object` type of class instances and, together with other optimizations, makes the encoded size of classes in the compact format much more competitive with that of structures.

Ice 3.5 and later use the compact format by default. Existing applications that rely on the slicing feature may begin receiving `NoObjectFactoryException` errors after migrating to Ice 3.5 or later, an indication that the run time does not recognize the most-derived type and cannot slice to a less-derived type because the compact format omits the necessary type information. These applications must use the sliced format instead, which behaves similarly to the format in version 1.0 of the encoding.

Using our previous example, an instance of `SubcontractorAccount` still includes the type ID string `"::TopLevelModule::IntermediateModule::SubcontractorAccount"` in the compact format, but omits the type ID strings for `Account` and `Object`. Clearly we could reduce this overhead even further by using shorter symbol names and eliminating the nested module, or we could omit the type ID string altogether by using a compact type ID.

Sliced Format

The sliced format in version 1.1 of the encoding retains the slicing behavior of version 1.0: a receiver discards the slices of unknown types until it finds a type that it recognizes, exhausts all available slices, or slices so far that the result would violate the formal signature of the operation. The sliced format still encodes the type IDs of the most-derived type as well as all intermediate types. However, unlike the 1.0 encoding, Ice no longer encodes the type ID for the base `Object` type in class instances such that, along with other improvements in the 1.1 encoding, the sliced format is still more efficient than the 1.0 encoding.

Copyright 2016, ZeroC, Inc.

Choosing a Format

Ice uses the compact format by default, so an application that needs slicing behavior must explicitly enable it as follows:

- Set the `Ice.Default.SlicedFormat` property to a non-zero value to force the Ice run time to use the sliced format by default.
- Annotate your Slice definitions with `format` metadata to selectively enable the sliced format for certain operations or interfaces.

For example, suppose an application can safely use the compact format most of the time, but still needs slicing in a few situations. In this case the application can use metadata to enable the sliced format where necessary:

Slice

```
interface Ledger
{
    Account getAccount(string id); // Uses compact format

    ["format:sliced"]
    Account importAccount(string source); // Uses sliced format
};
```

The semantics implied by these definitions state that the caller of `getAccount` assumes it will know every type that that might be returned, but the same cannot be said for `importAccount`. By enabling the sliced format here, we allow the client to "slice off" what it does not recognize, even if that means the client is left with only an instance of `Account` and not an instance of some derived type.

Now let's examine the opposite case: use the sliced format by default, and the compact format only in certain cases:

Slice

```
["format:sliced"]
interface Ledger
{
    ["format:compact"]
    Account getAccount(string id); // Uses compact format

    Account importAccount(string source); // Uses sliced format
};
```

Here we specify that all operations in `Ledger` use the sliced format unless overridden at the operation level, which we do for `getAccount`.

Keep the following points in mind when using the `format` metadata:

- The chosen format only controls how Ice encodes a value and has no effect on the decoding process.
- The format affects the input parameters, output parameters, and exceptions of an operation. Consider this example:

Slice

```
exception IncompatibleAccount { ... };

interface Ledger
{
    ["format:compact"]
    Account migrateAccount(Account oldAccount) throws IncompatibleAccount;
};
```


Copyright 2016, ZeroC, Inc.

The metadata forces the client to use the compact format when marshaling the input parameter `oldAccount`, and forces the server to use the compact format when marshaling the return value or the exception. For a given operation, it is not possible to use one format for the parameters and a different format for the exceptions.

- If you decide to use the `Ice.Default.SlicedFormat` property, be aware that this property only affects the sender of a class or exception. For example, if you enable this property in the client but not the server, then all objects sent by the client use the sliced format by default, but all objects and exceptions returned by the server use the compact format by default.

By offering two alternative formats, Ice gives you a great deal of flexibility in designing your applications. The compact format is ideal for applications that place a greater emphasis on efficiency, while the sliced format is helpful when clients and servers evolve independently.

Optional Objects

Declaring a data member of type `class` to be optional does not eliminate the Ice run time's requirement for type information. Suppose a client uses the following Slice definitions:

Slice

```
class UserInfo
{
    string name;
    optional(1) string organization;
};
```

Also suppose that the server is using a newer version of the definitions:

Slice

```
class GroupInfo
{
    ...;
};

class UserInfo
{
    string name;
    optional(1) string organization;
    optional(2) GroupInfo group;
};
```

When receiving an instance of `UserInfo`, the Ice run time in the client ignores the `group` member because it does not recognize its tag, but Ice must still be able to unmarshal (or skip) the `GroupInfo` instance data in the message. If the server used the compact format for its reply, the client will receive an exception while attempting to decode the message. On the other hand, this scenario would succeed when using the sliced format because the format embeds enough information for the receiver to skip instances of unknown types.

The lesson here is that adding a new class type to your application can have unexpected repercussions, especially when using the compact format. Remember that the *sender* can change formats without necessarily needing to update the receiver. For example, after adding the `group` member to `UserInfo`, the server could begin using the sliced format for all operations that return `UserInfo` to ensure that any existing clients would be able to ignore this member.

Preserving Slices

The concept of slicing involves discarding the slices of unknown types when decoding an instance of a Slice class or exception. Here is a

Copyright 2016, ZeroC, Inc.

simple example:

Slice

```
class Base
{
    int b;
};

class Intermediate extends Base
{
    int i;
};

class Derived extends Intermediate
{
    int d;
};

interface Relay
{
    Base transform(Base b);
};
```

The server implementing the `Relay` object must know the type `Base` (because it is statically referenced in the interface definition), but may not know `Intermediate` or `Derived`. Suppose the implementation of `transform` involves forwarding the object to another back-end server for processing and returning the transformed object to the caller. In effect, the `Relay` server is an intermediary. If the `Relay` server does not know the types `Intermediate` and `Derived`, it will slice an instance to `Base` and discard the data members of any more-derived types, which is clearly not the intended result because the back-end server *does* know those types. The only way the `Relay` server could successfully forward these objects is by knowing all possible derived types, which makes the application more difficult to evolve over time because the intermediary must be updated each time a new derived type is added.

To address this limitation, Ice supports a new metadata directive that allows an instance of a Slice class or exception to be forwarded with all of its slices intact, even if the intermediary does not know one or more of the instance's derived types. The new directive, `preserve-slice`, is shown below:

Copyright 2016, ZeroC, Inc.

Slice

```
["preserve-slice"]
class Base
{
    int b;
};

class Intermediate extends Base
{
    int i;
};

class Derived extends Intermediate
{
    int d;
};

["format:sliced"]
interface Relay
{
    Base transform(Base b);
};
```

With this change, all instances of `Base`, and types derived from `Base`, are encoded in a format that allows them to be forwarded intact. Also notice the addition of the `format:sliced` metadata on the `Relay` interface, which ensures that its operations use the sliced format and not the default compact format.

Unknown Sliced Objects

The combination of preserved slices and the sliced format means Ice applications now have a subtle new capability that was not available prior to Ice 3.5. Suppose we modify our `Relay` example as shown below:

Copyright 2016, ZeroC, Inc.

Slice

```
["preserve-slice"]
class Base
{
    int b;
};

class Intermediate extends Base
{
    int i;
};

class Derived extends Intermediate
{
    int d;
};

["format:sliced"]
interface Relay
{
    Object transform(Object b);
};
```

The only difference here is the signature of the `transform` operation, which now uses the `Object` type. Technically, it is not necessary for the intermediary server to know *any* of the class types that might be relayed via this new definition of `transform` because the formal types in its signature do not impose any requirements. As long as any object types known by the intermediary are marked with `preserve-slice`, and the `transform` operation uses the sliced format, this intermediary is capable of relaying objects of any type.

If the Ice run time in the intermediary does not know any of the types in an object's inheritance hierarchy, and the formal type is `Object`, Ice uses an instance of `UnknownSlicedObject` to represent the object:

C++

```
namespace Ice {
class UnknownSlicedObject : public Object, ... {
public:

    const std::string& getUnknownTypeId() const;
    SlicedDataPtr getSlicedData() const;

    ...
};
}
```

The implementation of `transform` receives an instance of `UnknownSlicedObject` and can use that object as its return value. If necessary, the implementation can determine the most-derived type of the object by calling `getUnknownTypeId`. The actual encoded state of the object is encapsulated by the `SlicedData` class:

Copyright 2016, ZeroC, Inc.

Slice

```
namespace Ice {
struct SliceInfo : public IceUtil::Shared {
    std::string typeId;
    std::vector<Byte> bytes;
    std::vector<ObjectPtr> objects;
    bool hasOptionalMembers;
    bool isLastSlice;
};


class SlicedData : public ... {
public:
    SlicedData(const SliceInfoSeq&);
    const SliceInfoSeq slices;

    ...
};
}
```

Applications do not normally need to use these types, but they provide all of the information the run time requires to forward an instance.

The definition of `UnknownSlicedObject`, `SlicedData` and `SliceInfo` are very similar in the other language mappings and are not shown here.

See Also

- Terminology
- Architectural Implications of Classes
- Classes Versus Structures
- Data Encoding for Classes
- Data Encoding for Exceptions
- Slice Metadata Directives
- Classes with Compact Type IDs

Copyright 2016, ZeroC, Inc.

Dynamic Ice

The Ice streaming API allows you to serialize and deserialize Slice types using either the Ice encoding or an encoding of your choice (for example, XML). This is useful, for example, if you want to store Slice types in a data base.

The dynamic invocation and dispatch interfaces allow you to write generic clients and servers that need not have compile-time knowledge of the Slice types used by an application. This makes it possible to create applications such as object browsers, protocol analyzers, or protocol bridges. In addition, the dynamic invocation and dispatch interfaces permit services such as IceStorm to be implemented without the need to unmarshal and remarshal every message, with considerable performance improvements.

Keep in mind that applications that use dynamic invocation and dispatch are tedious to implement and harder to prove correct (because what normally would be a compile-time error appears only as a run-time error with dynamic invocation and dispatch). Therefore, you should use the dynamic interfaces only if your application truly benefits from this trade-off.

Topics

- Streaming Interfaces
- Dynamic Invocation and Dispatch
- Asynchronous Dynamic Invocation and Dispatch

Copyright 2016, ZeroC, Inc.

Streaming Interfaces

Ice provides convenient interfaces for streaming Slice types to and from a sequence of bytes. You can use these interfaces in many situations, such as when serializing types for persistent storage, and when using Ice's dynamic invocation and dispatch facility.

The streaming interfaces are not defined in Slice, but are rather a collection of native classes provided by each language mapping.

The streaming interfaces are currently supported in C++, Java, and .NET.

A default implementation of the interfaces uses the Ice encoding, but other implementations are possible.

There are two primary abstract classes: `InputStream` and `OutputStream`. As you might guess, `InputStream` is used to extract Slice types from a sequence of bytes, while `OutputStream` is used to convert Slice types into a sequence of bytes. The classes provide the functions necessary to manipulate all of the core Slice types:

- Primitives (`bool`, `int`, `string`, etc.)
- Sequences of primitives
- Proxies
- Objects

The classes also provide functions that handle various details of the Ice encoding. Using these functions, you can manually insert and extract constructed types, such as dictionaries and structures, but doing so is tedious and error-prone. To make insertion and extraction of constructed types easier, the Slice compilers can optionally generate helper functions that manage the low-level details for you.

The remainder of this section describes the streaming interfaces for each supported language mapping. To properly use the streaming interfaces, you should be familiar with the Ice encoding. An example that demonstrates the use of the streaming interfaces is located in `demo/Ice/invoke` in your Ice distribution.

Topics

- C++ Streaming Interfaces
- Java Streaming Interfaces
- C-Sharp Streaming Interfaces

See Also

- Dynamic Invocation and Dispatch
- Data Encoding

Copyright 2016, ZeroC, Inc.

C++ Streaming Interfaces

We discuss the stream classes first, followed by the helper functions, and finish with an advanced use case.

Topics

- The InputStream Interface in C++
- The OutputStream Interface in C++
- Intercepting Object Insertion and Extraction in C++
- Intercepting User Exception Insertion and Extraction in C++

Copyright 2016, ZeroC, Inc.

The InputStream Interface in C++

On this page:

- The InputStream API in C++
- Extracting from an InputStream in C++
- Extracting Sequences of Built-In Types using Zero-Copy in C++
- Other InputStream Methods in C++

The `InputStream` API in C++

An `InputStream` that uses the Ice encoding can be created using the following functions:

C++

```
namespace Ice {
    InputStreamPtr createInputStream(
        const CommunicatorPtr& communicator,
        const std::vector<Ice::Byte>& data);

    InputStreamPtr createInputStream(
        const CommunicatorPtr& communicator,
        const std::vector<Ice::Byte>& data,
        const EncodingVersion& version);

    InputStreamPtr createInputStream(
        const CommunicatorPtr& communicator,
        const std::pair<const Ice::Byte*, const Ice::Byte*>& data);

    InputStreamPtr createInputStream(
        const CommunicatorPtr& communicator,
        const std::pair<const Ice::Byte*, const Ice::Byte*>& data,
        const EncodingVersion& version);

    InputStreamPtr wrapInputStream(
        const CommunicatorPtr& communicator,
        const std::vector<Ice::Byte>& data);

    InputStreamPtr wrapInputStream(
        const CommunicatorPtr& communicator,
        const std::vector<Ice::Byte>& data,
        const EncodingVersion& version);

    InputStreamPtr wrapInputStream(
        const CommunicatorPtr& communicator,
        const std::pair<const Ice::Byte*, const Ice::Byte*>& data);

    InputStreamPtr wrapInputStream(
        const CommunicatorPtr& communicator,
        const std::pair<const Ice::Byte*, const Ice::Byte*>& data,
        const EncodingVersion& version);
}
```

You can optionally specify an encoding version for the stream, otherwise the stream uses the communicator's default encoding version.

Copyright 2016, ZeroC, Inc.

Note that the `createInputStream` functions make a copy of the supplied data, whereas the `wrapInputStream` functions avoid copies by keeping a reference to the data. If you use `wrapInputStream`, it is your responsibility to ensure that the memory buffer remains valid and unmodified for the lifetime of the `InputStream` object.

The `InputStream` class is shown below.

C++

```
namespace Ice {
    class InputStream : ... {
    public:
        virtual CommunicatorPtr communicator() const = 0;

        virtual void sliceObjects(bool slice) = 0;

        virtual void read(bool& v) = 0;
        virtual void read(Byte& v) = 0;
        virtual void read(Short& v) = 0;
        virtual void read(Int& v) = 0;
        virtual void read(Long& v) = 0;
        virtual void read(Float& v) = 0;
        virtual void read(Double& v) = 0;
        virtual void read(std::string& s, bool convert = true) = 0;
        virtual void read(const char*& s, size_t& slen) = 0;
        virtual void read(const char*& s, size_t& slen, std::string& holder) = 0;
        virtual void read(std::vector<std::string>& s, bool convert) = 0;
        virtual void read(std::wstring& s) = 0;

        virtual void read(::std::vector<bool>&) = 0;

        Int readEnum(Int maxValue) { ... }

        template<typename T> inline void
        read(T& v) { ... }

        virtual void read(std::vector<std::string>& v, bool convert) = 0;

        virtual void read(std::pair<const bool*, const bool*>&,
                          IceUtil::ScopedArray<bool>&) = 0;

        virtual void read(std::pair<const Byte*, const Byte*>&) = 0;

        virtual void read(std::pair<const Short*, const Short*>&,
                          IceUtil::ScopedArray<Short>&) = 0;

        virtual void read(std::pair<const Int*, const Int*>&,
                          IceUtil::ScopedArray<Int>&) = 0;

        virtual void read(std::pair<const Long*, const Long*>&,
                          IceUtil::ScopedArray<Long>&) = 0;

        virtual void read(std::pair<const Float*, const Float*>&,
                          IceUtil::ScopedArray<Float>&) = 0;

        virtual void read(std::pair<const Double*, const Double*>&,
```


Copyright 2016, ZeroC, Inc.

```
                        IceUtil::ScopedArray<Double>&) = 0;

    virtual Int readSize() = 0;
    virtual Int readAndCheckSeqSize(int minWireSize) = 0;

    virtual ObjectPrx readProxy() = 0;

    template<typename T> inline void
    read(IceInternal::ProxyHandle<T>& v) { ... }

    virtual void readObject(const ReadObjectCallbackPtr& cb) = 0;

    template<typename T> inline void
    read(IceInternal::Handle<T>& v) { ... }

    virtual void throwException() = 0;
    virtual void throwException(const UserExceptionReaderFactoryPtr&) = 0;

    virtual void startObject() = 0;
    virtual SlicedDataPtr endObject(bool preserve) = 0;

    virtual void startException() = 0;
    virtual SlicedDataPtr endException(bool preserve) = 0;

    virtual std::string startSlice() = 0;
    virtual void endSlice() = 0;
    virtual void skipSlice() = 0;

    virtual EncodingVersion startEncapsulation() = 0;
    virtual void endEncapsulation() = 0;
    virtual EncodingVersion skipEncapsulation() = 0;

    virtual EncodingVersion getEncoding() = 0;

    virtual void readPendingObjects() = 0;

    virtual void rewind() = 0;

    virtual void skip(Int sz) = 0;
    virtual void skipSize() = 0;

    virtual bool readOptional(Int tag, OptionalFormat fmt) = 0;

    virtual void closure(void* p) = 0;
    virtual void* closure() const = 0;
};
```


Copyright 2016, ZeroC, Inc.

```
    typedef ... InputStreamPtr;
}
```

Extracting from an InputStream in C++

An `InputStream` provides a number of overloaded `read` member functions that allow you to read any parameter from the stream by simply calling `read`.

For example, you can extract a double value followed by a string from a stream as follows:

C++

```
vector<Ice::Byte> data = ...;
in = Ice::createInputStream(communicator, data);
double d;
in->read(d);
string s;
in->read(s);
```

Likewise, you can extract a sequence of a built-in type, or a complex type, or any other type from the stream as follows:

C++

```
vector<Ice::Byte> data = ...;
in = Ice::createInputStream(communicator, data);
// ...
IntSeq s; // Slice: sequence<int> IntSeq;
in->read(s);

ComplexType c;
in->read(c);
```

Extracting Sequences of Built-In Types using Zero-Copy in C++

`InputStream` provides a number of overloads that accept a pair of pointers. For example, you can extract a sequence of bytes as follows:

C++

```
vector<Ice::Byte> data = ...;
in = Ice::wrapInputStream(communicator, data);
std::pair<const Ice::Byte*, const Ice::Byte*> p;
in->read(p);
```

The same extraction works for the other built-in integral and floating-point types, such `int` and `double`.

If the extraction is for a byte sequence, the returned pointers always point at memory in the stream's internal marshaling buffer.

Copyright 2016, ZeroC, Inc.

For the other built-in types, the pointers refer to the internal marshaling buffer only if the Ice encoding is compatible with the machine and compiler representation of the type, otherwise the pointers refer to a temporary array allocated to hold the unmarshaled data. The overloads for zero-copy extraction accept an additional parameter of type `IceUtil::ScopedArray` that holds this temporary array when necessary.

Here is an example to illustrate how to extract a sequence of integers, regardless of whether the machine's encoding of integers matches the on-the-wire representation or not:

C++

```
#include <IceUtil/ScopedArray.h>
...
in = Ice::wrapInputStream(communicator, data);
std::pair<const Ice::Int*, const Ice::Int*> p;
IceUtil::ScopedArray<Ice::Int> a;
in->read(p, a);

for(const Ice::Int* i = p.first; i != p.second; ++i) {
    cout << *i << endl;
}
```

If the on-the-wire encoding matches that of the machine, and therefore zero-copy is possible, the returned pair of pointers points into the run time's internal marshaling buffer. Otherwise, the run time allocates an array, unmarshals the data into the array, and sets the pair of pointers to point into that array. Use of the `ScopedArray` helper template ensures that the array is deallocated once you let the `ScopedArray` go out of scope, so there is no need to call `delete[]`. `ScopedArray` is comparable to a `std::unique_ptr` for arrays.

Other `InputStream` Methods in C++

The remaining member functions of `InputStream` have the following semantics:

- `void sliceObjects(bool slice)`
  Determines the behavior of the stream when extracting Ice objects. An Ice object is "sliced" when a factory cannot be found for a Slice type ID, resulting in the creation of an object of a less-derived type. Slicing is typically disabled when the application expects all object factories to be present, in which case the exception `NoObjectFactoryException` is raised. The default behavior is to allow slicing.

- `void read(std::string& v, bool convert = true)`
  `void read(std::vector<std::string>& v, bool convert)`
  The bool argument determines whether the strings unmarshaled by these methods are processed by the string converter, if one is installed. The default behavior is to convert the strings.

- `void read(const char*& s, size_t& slen)`
  `void read(const char*& s, size_t& slen, std::string& holder)`
  The first read with only two arguments reads a string without string conversion, even if a string converter is installed; `s` points to memory managed by the `InputStream`. The second read with three arguments uses the string converter, if one is installed, and puts the resulting converted string in the `holder` parameter; `holder` is empty when no string conversion occurs.

- `Int readEnum(Int maxValue)`
  Unmarshals the integer value of an enumerator. The `maxValue` argument represents the highest enumerator value in the enumeration. Consider the following definitions:

  Slice

  ```
  enum Color { red, green, blue };
  enum Fruit { Apple, Pear=3, Orange };
  ```

  The maximum value for `Color` is 2, and the maximum value for `Fruit` is 4.

Copyright 2016, ZeroC, Inc.

In general, you should simply use `read` for your enum values. `read` with an enum parameter calls `readEnum` with the `maxValue` provided by the code generated by `slice2cpp`.

- `Ice::Int readSize()`
  The Ice encoding has a compact representation to indicate size. This function extracts a size and returns it as an integer.

- `Ice::Int readAndCheckSeqSize(int minWireSize)`
  Like `readSize`, this function reads a size and returns it, but also verifies that there is enough data remaining in the unmarshaling buffer to successfully unmarshal the elements of the sequence. The `minWireSize` parameter indicates the smallest possible on-the-wire representation of a single sequence element. If the unmarshaling buffer contains insufficient data to unmarshal the sequence, the function throws `UnmarshalOutOfBoundsException`.

- `Ice::ObjectPrx readProxy()`
  This function returns an instance of the base proxy type, `ObjectPrx`. Calling `read` with a proxy parameter has the same effect.

- `void readObject(const Ice::ReadObjectCallbackPtr &)`
  The Ice encoding for class instances requires extraction to occur in stages. The `readObject` function accepts a callback object of type `ReadObjectCallback`, whose definition is shown below:

  C++

```
namespace Ice {
    class ReadObjectCallback : ... {
    public:
        virtual void invoke(const Ice::ObjectPtr&) = 0;
    };
    typedef ... ReadObjectCallbackPtr;
}
```

  When the object instance is available, the callback object's `invoke` member function is called. The application must call `readPendingObjects` to ensure that all instances are properly extracted. If you are not interested in receiving a callback when the object is extracted, it is easier to call `read` with a `Ptr` parameter. Note that calling `endEncapsulation` implicitly calls `readPendingObjects` if necessary.

- `void throwException()`
  `void throwException(const UserExceptionReaderFactoryPtr& factory)`
  These functions extract a user exception from the stream and throw it. If the stored exception is of an unknown type, the functions attempt to extract and throw a less-derived exception. If that also fails, an exception is thrown: for the 1.0 encoding, the exception is `UnmarshalOutOfBoundsException`, for the 1.1 encoding, the exception is `UnknownUserException`.

- `void startObject()`
  `SlicedDataPtr endObject(bool preserve)`
  The `startObject` method must be called prior to reading the slices of an object. The `endObject` method must be called after all slices have been read. Pass true to `endObject` in order to preserve the slices of any unknown more-derived types, or false to discard the slices. If `preserve` is true and the stream actually preserved any slices, the return value of `endObject` is a non-nil `SlicedData` object that encapsulates the slice data. If the caller later wishes to forward the object with any preserved slices intact, it must supply this `SlicedData` object to the output stream.

- `void startException()`
  `SlicedDataPtr endException(bool preserve)`
  The `startException` method must be called prior to reading the slices of an exception. The `endException` method must be called after all slices have been read. Pass true to `endException` in order to preserve the slices of any unknown more-derived types, or false to discard the slices. If `preserve` is true and the stream actually preserved any slices, the return value of `endException` is a non-nil `SlicedData` object that encapsulates the slice data. If the caller later wishes to forward the exception with any preserved slices intact, it must supply this `SlicedData` object to the output stream.

- `std::string startSlice()`
  `void endSlice()`
  `void skipSlice()`
  Start, end, and skip a slice of member data, respectively. These functions are used when manually extracting the slices of an Ice object or user exception. The `startSlice` method returns the type ID of the next slice, which may be an empty string depending on the format used to encode the object or exception.

Copyright 2016, ZeroC, Inc.

- `EncodingVersion startEncapsulation()`
  `void endEncapsulation()`
  `EncodingVersion skipEncapsulation()`
  Start, end, and skip an encapsulation, respectively. The `startEncapsulation` and `skipEncapsulation` methods return the encoding version used to encode the contents of the encapsulation.

- `EncodingVersion getEncoding()`
  Returns the encoding version currently in use by the stream.

- `void readPendingObjects()`
  An application must call this function after all other data has been extracted. This function extracts the state of Ice objects and invokes their corresponding callback objects (see `readObject`). For backward compatibility with encoding version 1.0, this function must only be called when non-optional data members or parameters use class types.

- `void rewind()`
  Resets the position of the stream to the beginning.

- `void skip(Int sz)`
  Skips the given number of bytes.

- `void skipSize()`
  Reads a size at the current position and skips that number of bytes.

- `bool readOptional(Int tag, OptionalFormat fmt)`
  Returns true if an optional value with the given tag and format is present, or false otherwise. If this method returns true, the data associated with that optional value must be read next. Optional values must be read in order by tag from least to greatest. The `OptionalFormat` enumeration is defined as follows:

  C++

  ```
  namespace Ice {
      enum OptionalFormat {
          OptionalFormatF1, OptionalFormatF2, OptionalFormatF4,
  OptionalFormatF8,
          OptionalFormatSize, OptionalFormatVSize, OptionalFormatFSize,
          OptionalFormatEndMarker
      };
  }
  ```

  Refer to the encoding discussion for more information on the meaning of these values.

- `void closure(void* p)`
  `void* closure() const`
  Allows an arbitrary value to be associated with the stream.

See Also

- Smart Pointers for Classes
- slice2cpp Command-Line Options
- C++ Strings and Character Encoding
- Data Encoding for Classes
- Basic Data Encoding
- The C++ ScopedArray Template

Copyright 2016, ZeroC, Inc.

The OutputStream Interface in C++

On this page:

- The OutputStream API in C++
- Inserting into an OutputStream in C++
- Inserting Sequences of Built-In Types in C++
- Other OutputStream Member Functions in C++

The `OutputStream` API in C++

An `OutputStream` is created using the following functions:

C++

```
namespace Ice {
    OutputStreamPtr createOutputStream(const CommunicatorPtr& communicator);

    OutputStreamPtr createOutputStream(constCommunicatorPtr& communicator,
                                       const EncodingVersion& version);
}
```

You can optionally specify an encoding version for the stream, otherwise the stream uses the communicator's default encoding version.

The `OutputStream` class is shown below.

C++

```
namespace Ice {
    class OutputStream : ... {
    public:
        virtual CommunicatorPtr communicator() const = 0;

        virtual void write(bool) = 0;
        virtual void write(Byte) = 0;
        virtual void write(Short) = 0;
        virtual void write(Int) = 0;
        virtual void write(Long) = 0;
        virtual void write(Float) = 0;
        virtual void write(Double) = 0;
        virtual void write(const std::string&, bool = true) = 0;
        virtual void write(const char*, size_t, bool = true) = 0;
        virtual void write(const char*, bool = true) = 0;
        virtual void write(const std::wstring&) = 0;

        virtual void write(const std::vector<bool>&) = 0;

        virtual void write(const bool* begin, const bool* end) = 0;
        virtual void write(const Byte* begin, const Byte* end) = 0;
        virtual void write(const Short* begin, const Short* end) = 0;
        virtual void write(const Int* begin, const Int* end) = 0;
        virtual void write(const Long* begin, const Long* end) = 0;
        virtual void write(const Float* begin, const Float* end) = 0;
        virtual void write(const Double* begin, const Double* end) = 0;
```


Copyright 2016, ZeroC, Inc.

```
        virtual void write(const std::vector<std::string>& v, bool convert) = 0;

        void writeEnum(Int v, Int maxValue) { ... }

        virtual bool writeOptional(Int tag, OptionalFormat fmt) = 0;

        template<typename T> inline void
        write(const T& v) { ... }

        virtual void writeSize(Int sz) = 0;
        virtual void startSize() = 0;
        virtual void endSize() = 0;

        virtual void writeProxy(const ObjectPrx& v) = 0;

        template<typename T> inline void
        write(const IceInternal::ProxyHandle<T>& v) { ... }

        virtual void writeObject(const ObjectPtr&) = 0;

        template<typename T> inline void
        write(const IceInternal::Handle<T>& v) { ... }

        virtual void writeException(const UserException& e) = 0;

        virtual void startObject(const SlicedDataPtr& sd) = 0;
        virtual void endObject() = 0;

        virtual void startException(const SlicedDataPtr& sd) = 0;
        virtual void endException() = 0;

        virtual void startSlice(const std::string& typeId, bool last) = 0;
        virtual void endSlice() = 0;

        virtual void startEncapsulation(const EncodingVersion& v, FormatType fmt) =
0;
        virtual void startEncapsulation() = 0;
        virtual void endEncapsulation() = 0;

        virtual EncodingVersion getEncoding() const = 0;

        virtual void writePendingObjects() = 0;

        virtual void finished(std::vector<Byte>& v) = 0;
        virtual std::pair<const Byte*, const Byte*> finished() = 0;

        virtual size_type pos() = 0;
        virtual void rewrite(Int v, size_type pos) = 0;

        virtual void reset(bool) = 0;
```


Copyright 2016, ZeroC, Inc.

```
    };
}
```

Inserting into an OutputStream in C++

An `OutputStream` provides a number of overloaded `write` member functions that allow you to insert any parameter into the stream by simply calling `write`.

For example, you can insert a double value followed by a string into a stream as follows:

C++

```
out = Ice::createOutputStream(communicator);
Ice::Double d = 3.14;
out->write(d);
string s = "Hello";
out->write(s);
```

Likewise, you can insert a sequence of built-in type, or a sequence of a complex constructed type, or any other type, with the same syntax:

C++

```
out = Ice::createOutputStream(communicator);
IntSeq s = ...;
out->write(s);

ComplexType c = ...;
out->write(c);
```

Inserting Sequences of Built-In Types in C++

`OutputStream` provides a number of overloads that accept a pair of pointers. For example, you can insert a sequence of bytes as follows:

C++

```
out = Ice::createOutputStream(communicator);
vector<Ice::Byte> v = ...;
out->write(&v[0], &v[v.size()]);
```

The same insertion technique works for the other built-in integral and floating-point types, such `int` and `double`. Insertion with these functions is a straight copy into the marshaling buffer when the internal representation of the data in memory is the same as the on-the-wire representation. (Note that the two pointers must point at a contiguous block of memory.)

Other `OutputStream` Member Functions in C++

The remaining member functions of `OutputStream` have the following semantics:

- `void write(const std::string& v, bool convert = true)`
  `void write(const char* v, size_t vlen, bool convert = true)`

Copyright 2016, ZeroC, Inc.

```
void write(const char* v, bool convert = true)
void write(const std::vector<std::string>&, bool convert)
```
The boolean argument determines whether the strings marshaled by these methods are processed by the string converter, if one is installed. The default behavior is to convert the strings.

- `void writeEnum(Int val, Int maxValue)`
Writes the integer value of an enumerator. The `maxValue` argument represents the highest enumerator value in the enumeration. Consider the following definitions:

Slice

```
enum Color { red, green, blue };
enum Fruit { Apple, Pear=3, Orange };
```

The maximum value for `Color` is 2, and the maximum value for `Fruit` is 4.
In general, you should simply use `write` for your enum values. `write` with an enum parameter calls `writeEnum` with the `maxValue` provided by the code generated by `slice2cpp`.

- `void writeSize(Ice::Int sz)`
The Ice encoding has a compact representation to indicate size. This function converts the given non-negative integer into the proper encoded representation.

- `void startSize()`
`void endSize()`
The encoding for optional values uses a 32-bit integer to hold the size of variable-length types. Calling `startSize` writes a placeholder value for the size; after writing the data, call `endSize` to patch the placeholder with the actual size.

- `bool writeOptional(Int tag, OptionalFormat fmt)`
Prepares the stream to write an optional value with the given tag and format. Returns true if the value should be written, or false otherwise. A return value of false indicates that the encoding version in use by the stream does not support optional values. If this method returns true, the data associated with that optional value must be written next. Optional values must be written in order by tag from least to greatest. The `OptionalFormat` enumeration is defined as follows:

C++

```
namespace Ice {
    enum OptionalFormat {
        OptionalFormatF1, OptionalFormatF2, OptionalFormatF4,
OptionalFormatF8,
        OptionalFormatSize, OptionalFormatVSize, OptionalFormatFSize,
        OptionalFormatEndMarker
    };
}
```

Refer to the encoding discussion for more information on the meaning of these values.

- `void writeProxy(const Ice::ObjectPrx&)`
Inserts a proxy into the stream. It is equivalent to calling `write` with a proxy parameter.

- `void writeObject(const Ice::ObjectPtr&)`
Inserts an Ice object. The Ice encoding for class instances may cause the insertion of this object to be delayed, in which case the stream retains a reference to the given object and the stream does not insert its state it until `writePendingObjects` is invoked on the stream. It is equivalent to calling `write` with a `Ptr` parameter.

- `void writeException(const Ice::UserException & ex)`
Inserts a user exception. It is equivalent to calling `write` with a user exception. The exception argument may be an instance of `UserExceptionWriter`, which allows you to implement your own exception marshaling logic.

- `void startObject(const SlicedDataPtr& sd)`
`void endObject()`

Copyright 2016, ZeroC, Inc.

When marshaling the slices of an object, the application must first call `startObject`, then marshal the slices, and finally call `endObject`. The caller can pass a `SlicedData` object containing the preserved slices of unknown more-derived types, or 0 if there are no preserved slices.

- `void startException(const SlicedDataPtr& sd)`
  `void endException()`
  When marshaling the slices of an exception, the application must first call `startException`, then marshal the slices, and finally call `endException`. The caller can pass a `SlicedData` object containing the preserved slices of unknown more-derived types, or 0 if there are no preserved slices.

- `void startSlice(const std::string& typeId, bool last)`
  `void endSlice()`
  Starts and ends a slice of object or exception member data. The call to `startSlice` must include the type ID for the current slice, and a boolean indicating whether this is the last slice of the object or exception.

- `void startEncapsulation(const EncodingVersion& v, FormatType fmt)`
  `void startEncapsulation()`
  `void endEncapsulation()`
  Starts and ends an encapsulation, respectively. The first overloading of `startEncapsulation` allows you to specify the encoding version as well as the format to use for any objects and exceptions marshaled within this encapsulation.

- `EncodingVersion getEncoding()`
  Returns the encoding version currently being used by the stream.

- `void writePendingObjects()`
  Encodes the state of Ice objects whose insertion was delayed during `writeObject`. This member function must only be called once. For backward compatibility with encoding version 1.0, this function must only be called when non-optional data members or parameters use class types.

- `void finished(std::vector< Ice::Byte >& data)`
  `std::pair<const Byte*, const Byte*> finished()`
  Indicates that marshaling is complete. This member function must only be called once. In the first overloading, the given byte sequence is filled with a copy of the encoded data. The second overloading avoids a copy by returning a pair of pointers to the stream's internal memory; these pointers are valid for the lifetime of the `OutputStream` object.

- `size_type pos()`
  `void rewrite(Int v, size_type pos)`
  The `pos` method returns the stream's current position, and `rewrite` allows you to overwrite a 32-bit integer value at the given position in the stream. Calling `rewrite` does not change the stream's current position.

- `void reset(bool clearBuffer)`
  Resets the writing position of the stream to the beginning. If `clearBuffer` is true, the stream releases the memory it has allocated to hold the encoded data.

See Also

- Basic Data Encoding
- slice2cpp Command-Line Options
- C++ Strings and Character Encoding
- Data Encoding for Classes
- Data Encoding for Exceptions

Copyright 2016, ZeroC, Inc.

Intercepting Object Insertion and Extraction in C++

In some situations it may be necessary for an application to intercept the insertion and extraction of Ice objects. For example, the Ice extension for PHP is implemented using Ice for C++ but represents Ice objects as native PHP objects. The PHP extension accomplishes this by manually encoding and decoding Ice objects as directed by the data encoding rules. However, the extension obviously cannot pass a native PHP object to the C++ stream function for writing objects. To bridge this gap between object systems, Ice supplies the helper classes described below.

On this page:

- Inserting Objects in C++
- Extracting Objects in C++

Inserting Objects in C++

The `ObjectWriter` base class facilitates the insertion of objects to a stream:

C++

```
namespace Ice {
    class ObjectWriter : public Ice::Object {
    public:
        virtual void write(const OutputStreamPtr&) const = 0;
        // ...
    };
    typedef ... ObjectWriterPtr;
}
```

A foreign Ice object is inserted into a stream using the following technique:

1. A C++ "wrapper" class is derived from `ObjectWriter`. This class wraps the foreign object and implements the `write` member function.
2. An instance of the wrapper class is passed to `writeObject`. (This is possible because `ObjectWriter` derives from `Ice::Object`.) Eventually, the `write` member function is invoked on the wrapper instance.
3. The implementation of `write` encodes the object's state as directed by the data encoding rules for classes.

It is the application's responsibility to ensure that there is a one-to-one mapping between foreign Ice objects and wrapper objects. This is necessary in order to ensure the proper encoding of object graphs.

Extracting Objects in C++

The `ObjectReader` class facilitates the extraction of objects from a stream:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {
    class ObjectReader : public Ice::Object {
    public:
        virtual void read(const InputStreamPtr&) = 0;
        // ...
    };
    typedef ... ObjectReaderPtr;

    class CompactIdResolver : public IceUtil::Shared {
    public:
        virtual std::string resolve(Ice::Int) const = 0;
    };
    typedef ... CompactIdResolverPtr;
}
```

Extracting the state of a foreign Ice object is more complicated than insertion:

1. A C++ "wrapper" class is derived from `ObjectReader`. An instance of this class represents a foreign Ice object.
2. An object factory is installed that returns instances of the wrapper class. Note that a single object factory can be used for all Slice types if it is registered with an empty Slice type ID.
3. A C++ callback class is derived from `ReadObjectCallback`. The implementation of `invoke` expects its argument to be either nil or an instance of the wrapper class as returned by the object factory.
4. An instance of the callback class is passed to `readObject`.
5. When the stream is ready to extract the state of an object, it invokes `read` on the wrapper class. The implementation of `read` decodes the object's state as directed by the data encoding rules for classes.
6. The callback object passed to `readObject` is invoked, passing the instance of the wrapper object. All other callback objects representing the same instance in the stream (in case of object graphs) are invoked with the same wrapper object.

If your class definitions use compact type IDs, you must also supply an implementation of `CompactIdResolver` when initializing the communicator. This object is responsible for translating numeric type IDs into their string equivalents.

See Also

- Client-Side Slice-to-PHP Mapping
- Data Encoding
- Data Encoding for Classes
- C++ Mapping for Classes
- Communicator Initialization

Copyright 2016, ZeroC, Inc.

Intercepting User Exception Insertion and Extraction in C++

Inserting a User Exception in C++

As in the case of Ice objects, a Dynamic Ice application may represent user exceptions in a native format that is not directly compatible with the Ice API. If the application needs to raise such a user exception to the Ice run time (so that it will be marshaled and sent back to the client), the exception must be wrapped in a subclass of `Ice::UserException`. The Dynamic Ice API provides a class to simplify this process:

C++

```
namespace Ice {
    class UserExceptionWriter : public UserException {
    public:
        UserExceptionWriter(const CommunicatorPtr&);

        virtual void write(const OutputStreamPtr&) const = 0;

        virtual std::string ice_name() const = 0;
        virtual Ice::Exception* ice_clone() const = 0;
        virtual void ice_throw() const = 0;

        // ...
    };
    typedef ... UserExceptionWriterPtr;
}
```

A subclass of `UserExceptionWriter` is responsible for supplying a communicator to the constructor, and for implementing the following methods:

- `void write(const OutputStreamPtr&) const`
  This method is invoked when the Ice run time is ready to marshal the exception. The subclass must marshal the exception using the encoding rules for exceptions.
- `std::string ice_name() const`
  Return the Slice name of the exception.
- `Ice::Exception* ice_clone() const`
  Return a copy of the exception.
- `void ice_throw() const`
  Raise the exception by calling `throw *this`.

Extracting a User Exception in C++

An application extracts a user exception by calling one of two versions of the `throwException` method defined in the `InputStream` class:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {
    class InputStream : ... {
    public:
        virtual void throwException() = 0;
        virtual void throwException(const UserExceptionReaderFactoryPtr&) = 0;

        // ...
    };
}
```

The version without any arguments attempts to locate and throw a C++ implementation of the encoded exception, relying on statically-generated type information emitted by the Slice-to-C++ compiler.

If your goal is to create an exception in another type system, such as a native PHP exception object, you must call the second version of `throwException` and pass an implementation of `UserExceptionReaderFactory`:

C++

```
namespace Ice {
    class UserExceptionReaderFactory : public IceUtil::Shared {
    public:
        virtual void createAndThrow(const std::string& typeId) const = 0;
    };
    typedef ... UserExceptionReaderFactoryPtr;
}
```

As the stream iterates over slices of an exception from most-derived to least-derived, it invokes `createAndThrow` passing the type ID of each slice, giving the application an opportunity to raise an instance of `UserExceptionReader`:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {
    class UserExceptionReader : public UserException {
    public:
        UserExceptionReader(const CommunicatorPtr&);
        ~UserExceptionReader() throw();

        virtual void read(const InputStreamPtr&) const = 0;

        virtual std::string ice_name() const = 0;
        virtual UserException* ice_clone() const = 0;
        virtual void ice_throw() const = 0;
    protected:
        const CommunicatorPtr _communicator;
    };
}
```

Subclasses of `UserExceptionReader` must implement the pure virtual functions. In particular, the implementation of `read` must call `InputStream::startException`, unmarshal the remaining slices, and then call `InputStream::endException`.

See Also

- Intercepting Object Insertion and Extraction in C++
- Data Encoding for Exceptions

Copyright 2016, ZeroC, Inc.

Java Streaming Interfaces

We discuss the stream classes first, followed by the helper functions, and finish with an advanced use case.

Topics

- The InputStream Interface in Java
- The OutputStream Interface in Java
- Intercepting Object Insertion and Extraction in Java
- Intercepting User Exception Insertion and Extraction in Java
- Stream Helper Methods in Java

Copyright 2016, ZeroC, Inc.

The InputStream Interface in Java

An `InputStream` is created using the following functions:

Java

```
package Ice;

public class Util {
    public static InputStream
    createInputStream(Communicator communicator, byte[] data);

    public static InputStream
    createInputStream(Communicator communicator, byte[] data, EncodingVersion
version);

    public static InputStream
    wrapInputStream(Communicator communicator, byte[] data);

    public static InputStream
    wrapInputStream(Communicator communicator, byte[] data, EncodingVersion
version);
}
```

You can optionally specify an encoding version for the stream, otherwise the stream uses the communicator's default encoding version.

Note that the `createInputStream` functions make a copy of the supplied data, whereas the `wrapInputStream` functions avoid copies by keeping a reference to the data. If you use `wrapInputStream`, it is your responsibility to ensure that the byte array remains unmodified for the lifetime of the `InputStream` object.

The `InputStream` interface is shown below.

Java

```
package Ice;

public interface InputStream {
    Communicator communicator();

    void sliceObjects(boolean slice);

    boolean readBool();
    boolean[] readBoolSeq();

    byte readByte();
    byte[] readByteSeq();

    short readShort();
    short[] readShortSeq();

    int readInt();
    int[] readIntSeq();
```


Copyright 2016, ZeroC, Inc.

```
long readLong();
long[] readLongSeq();

float readFloat();
float[] readFloatSeq();

double readDouble();
double[] readDoubleSeq();

String readString();
String[] readStringSeq();

int readSize();
int readAndCheckSeqSize(int minSizeWireSize);

ObjectPrx readProxy();

void readObject(ReadObjectCallback cb);

int readEnum(int maxValue);

void throwException() throws UserException;
void throwException(UserExceptionReaderFactory factory) throws UserException;

void startObject();
SlicedData endObject(boolean preserve);

void startException();
SlicedData endException(boolean preserve);

String startSlice();
void endSlice();
void skipSlice();

EncodingVersion startEncapsulation();
void endEncapsulation();
EncodingVersion skipEncapsulation();

EncodingVersion getEncoding();

void readPendingObjects();
java.io.Serializable readSerializable();

void rewind();

void skip(int sz);
void skipSize();

boolean readOptional(int tag, OptionalFormat format);

int pos();
```


Copyright 2016, ZeroC, Inc.

```
    void destroy();
}
```

Member functions are provided for extracting all of the primitive types, as well as sequences of primitive types; these are self-explanatory. The remaining member functions have the following semantics:

- `void sliceObjects(boolean slice)` Determines the behavior of the stream when extracting Ice objects. An Ice object is "sliced" when a factory cannot be found for a Slice type ID, resulting in the creation of an object of a less-derived type. Slicing is typically disabled when the application expects all object factories to be present, in which case the exception `NoObjectFactoryException` is raised. The default behavior is to allow slicing.
- `int readSize()`
  The Ice encoding has a compact representation to indicate size. This function extracts a size and returns it as an integer.
- `int readAndCheckSeqSize(int minWireSize)`
  Like `readSize`, this function reads a size and returns it, but also verifies that there is enough data remaining in the unmarshaling buffer to successfully unmarshal the elements of the sequence. The `minWireSize` parameter indicates the smallest possible on-the-wire representation of a single sequence element. If the unmarshaling buffer contains insufficient data to unmarshal the sequence, the function throws `UnmarshalOutOfBoundsException`.
- `Ice.ObjectPrx readProxy()`
  This function returns an instance of the base proxy type, `ObjectPrx`. The Slice compiler optionally generates helper functions to extract proxies of user-defined types.
- `void readObject(ReadObjectCallback cb)`
  The Ice encoding for class instances requires extraction to occur in stages. The `readObject` function accepts a callback object of type `ReadObjectCallback`, whose definition is shown below:

  Java

  ```
  package Ice;

  public interface ReadObjectCallback {
      void invoke(Ice.Object obj);
  }
  ```

  When the object instance is available, the callback object's `invoke` member function is called. The application must call `readPendingObjects` to ensure that all instances are properly extracted. Note that applications rarely need to invoke this member function directly; the helper functions generated by the Slice compiler are easier to use.
- `int readEnum(int maxValue)`
  Unmarshals the integer value of an enumerator. The `maxValue` argument represents the highest enumerator value in the enumeration. Consider the following definitions:

  Slice

  ```
  enum Color { red, green, blue };
  enum Fruit { Apple, Pear=3, Orange };
  ```

  The maximum value for `Color` is 2, and the maximum value for `Fruit` is 4.
- `void throwException() throws UserException`

  `void throwException(UserExceptionReaderFactory factory)`

  These functions extract a user exception from the stream and throw it. If the stored exception is of an unknown type, the functions

Copyright 2016, ZeroC, Inc.

attempt to extract and throw a less-derived exception. If that also fails, an exception is thrown: for the 1.0 encoding, the exception is `UnmarshalOutOfBoundsException`, for the 1.1 encoding, the exception is `UnknownUserException`.

- `String startSlice()`
  `void endSlice()`
  `void skipSlice()`
  Start, end, and skip a slice of member data, respectively. These functions are used when manually extracting the slices of an Ice object or user exception. The `startSlice` method returns the type ID of the next slice, which may be an empty string depending on the format used to encode the object or exception.

- `void startObject()`
  `SlicedData endObject(boolean preserve)`
  The `startObject` method must be called prior to reading the slices of an object. The `endObject` method must be called after all slices have been read. Pass true to `endObject` in order to preserve the slices of any unknown more-derived types, or false to discard the slices. If `preserve` is true and the stream actually preserved any slices, the return value of `endObject` is a non-nil `SlicedData` object that encapsulates the slice data. If the caller later wishes to forward the object with any preserved slices intact, it must supply this `SlicedData` object to the output stream.

- `void startException()`
  `SlicedData endException(boolean preserve)`
  The `startException` method must be called prior to reading the slices of an exception. The `endException` method must be called after all slices have been read. Pass true to `endException` in order to preserve the slices of any unknown more-derived types, or false to discard the slices. If `preserve` is true and the stream actually preserved any slices, the return value of `endException` is a non-nil `SlicedData` object that encapsulates the slice data. If the caller later wishes to forward the exception with any preserved slices intact, it must supply this `SlicedData` object to the output stream.

- `EncodingVersion startEncapsulation()`
  `void endEncapsulation()`
  `EncodingVersion skipEncapsulation()`
  Start, end, and skip an encapsulation, respectively. The `startEncapsulation` and `skipEncapsulation` methods return the encoding version used to encode the contents of the encapsulation.

- `EncodingVersion getEncoding()`
  Returns the encoding version currently in use by the stream.

- `void readPendingObjects()`
  An application must call this function after all other data has been extracted, but only if Ice objects were encoded. This function extracts the state of Ice objects and invokes their corresponding callback objects (see `readObject`). For backward compatibility with encoding version 1.0, this function must only be called when non-optional data members or parameters use class types.

- `java.io.Serializable readSerializable()`
  Reads a serializable Java object from the stream.

- `void rewind()`
  Resets the position of the stream to the beginning.

- `void skip(int sz)`
  Skips the given number of bytes.

- `void skipSize()`
  Reads a size at the current position and skips that number of bytes.

- `boolean readOptional(int tag, OptionalFormat fmt)`
  Returns true if an optional value with the given tag and format is present, or false otherwise. If this method returns true, the data associated with that optional value must be read next. Optional values must be read in order by tag from least to greatest. The `OptionalFormat` enumeration is defined as follows:

Copyright 2016, ZeroC, Inc.

Java

```
package Ice;
enum OptionalFormat {
    OptionalFormatF1, OptionalFormatF2, OptionalFormatF4, OptionalFormatF8,
    OptionalFormatSize, OptionalFormatVSize, OptionalFormatFSize,
    OptionalFormatEndMarker
}
```

Refer to the encoding discussion for more information on the meaning of these values.

- `void destroy()`
  Applications must call this function in order to reclaim resources.

Here is a simple example that demonstrates how to extract a boolean and a sequence of strings from a stream:

Java

```
byte[] data = ...
Ice.InputStream in =
    Ice.Util.createInputStream(communicator, data);
try {
    boolean b = in.readBool();
    String[] seq = in.readStringSeq();
} finally {
    in.destroy();
}
```

See Also

- Basic Data Encoding
- Data Encoding for Classes
- Data Encoding for Exceptions
- Serializable Objects in Java
- Stream Helper Methods in Java

Copyright 2016, ZeroC, Inc.

The OutputStream Interface in Java

An `OutputStream` is created using the following function:

Java

```
package Ice;

public class Util {
    public static OutputStream
    createOutputStream(Communicator communicator);

    public static OutputStream
    createOutputStream(Communicator communicator, EncodingVersion version);
}
```

You can optionally specify an encoding version for the stream, otherwise the stream uses the communicator's default encoding version.

The `OutputStream` class is shown below.

Java

```
package Ice;

public interface OutputStream {
    Communicator communicator();

    void writeBool(boolean v);
    void writeBoolSeq(boolean[] v);

    void writeByte(byte v);
    void writeByteSeq(byte[] v);

    void writeShort(short v);
    void writeShortSeq(short[] v);

    void writeInt(int v);
    void writeIntSeq(int[] v);

    void writeLong(long v);
    void writeLongSeq(long[] v);

    void writeFloat(float v);
    void writeFloatSeq(float[] v);

    void writeDouble(double v);
    void writeDoubleSeq(double[] v);

    void writeString(String v);
    void writeStringSeq(String[] v);

    void writeSize(int sz);
```


Copyright 2016, ZeroC, Inc.

```
void writeProxy(ObjectPrx v);

void writeObject(Ice.Object v);

void writeEnum(int v, int maxValue);

void writeException(UserException ex);

void startObject(SlicedData sd);
void endObject();

void startException(SlicedData sd);
void endException();

void startSlice(String typeId, boolean last);
void endSlice();

void startEncapsulation(EncodingVersion encoding, FormatType format);
void startEncapsulation();
void endEncapsulation();

EncodingVersion getEncoding();

void writePendingObjects();

boolean writeOptional(int tag, OptionalFormat format);

int pos();

void rewrite(int sz, int pos);

void startSize();
void endSize();

byte[] finished();

void reset(boolean clearBuffer);

void writeSerializable(java.io.Serializable o);
```


Copyright 2016, ZeroC, Inc.

```
    void destroy();
}
```

Member functions are provided for inserting all of the primitive types, as well as sequences of primitive types; these are self-explanatory. The remaining member functions have the following semantics:

- `void writeSize(int sz)`
  The Ice encoding has a compact representation to indicate size. This function converts the given non-negative integer into the proper encoded representation.
- `void writeProxy(Ice.ObjectPrx v)`
  Inserts a proxy.
- `void writeObject(Ice.Object v)`
  Inserts an Ice object. The Ice encoding for class instances may cause the insertion of this object to be delayed, in which case the stream retains a reference to the given object and does not insert its state it until `writePendingObjects` is invoked on the stream.
- `void writeEnum(int val, int maxValue)`
  Writes the integer value of an enumerator. The `maxValue` argument represents the highest enumerator value in the enumeration. Consider the following definitions:

  Slice

  ```
  enum Color { red, green, blue };
  enum Fruit { Apple, Pear=3, Orange };
  ```

  The maximum value for `Color` is 2, and the maximum value for `Fruit` is 4.
- `void writeException(UserException ex)`
  Inserts a user exception. The exception argument may be an instance of `UserExceptionWriter`, which allows you to implement your own exception marshaling logic.
- `void startObject(SlicedData sd)`
  `void endObject()`
  When marshaling the slices of an object, the application must first call `startObject`, then marshal the slices, and finally call `endObject`. The caller can pass a `SlicedData` object containing the preserved slices of unknown more-derived types, or 0 if there are no preserved slices.
- `void startException(SlicedData sd)`
  `void endException()`
  When marshaling the slices of an exception, the application must first call `startException`, then marshal the slices, and finally call `endException`. The caller can pass a `SlicedData` object containing the preserved slices of unknown more-derived types, or 0 if there are no preserved slices.
- `void startSlice(String typeId, boolean last)`
  `void endSlice()`
  Starts and ends a slice of object or exception member data. The call to `startSlice` must include the type ID for the current slice, and a boolean indicating whether this is the last slice of the object or exception.
- `void startEncapsulation(EncodingVersion encoding, FormatType format)`
  `void startEncapsulation()`
  `void endEncapsulation()`
  Starts and ends an encapsulation, respectively. The first overloading of `startEncapsulation` allows you to specify the encoding version as well as the format to use for any objects and exceptions marshaled within this encapsulation.
- `EncodingVersion getEncoding()`
  Returns the encoding version currently being used by the stream.
- `void writePendingObjects()`
  Encodes the state of Ice objects whose insertion was delayed during `writeObject`. This member function must only be called once. For backward compatibility with encoding version 1.0, this function must only be called when non-optional data members or

Copyright 2016, ZeroC, Inc.

parameters use class types.

- `boolean writeOptional(int tag, OptionalFormat fmt)`
  Prepares the stream to write an optional value with the given tag and format. Returns true if the value should be written, or false otherwise. A return value of false indicates that the encoding version in use by the stream does not support optional values. If this method returns true, the data associated with that optional value must be written next. Optional values must be written in order by tag from least to greatest. The `OptionalFormat` enumeration is defined as follows:

  Java

  ```
  package Ice;
  enum OptionalFormat {
      OptionalFormatF1, OptionalFormatF2, OptionalFormatF4, OptionalFormatF8,
      OptionalFormatSize, OptionalFormatVSize, OptionalFormatFSize,
      OptionalFormatEndMarker
  }
  ```

  Refer to the encoding discussion for more information on the meaning of these values.

- `int pos()`
  `void rewrite(int sz, int pos)`
  The `pos` method returns the stream's current position, and `rewrite` allows you to overwrite a 32-bit integer value at the given position in the stream. Calling `rewrite` does not change the stream's current position.

- `void startSize()`
  `void endSize()`
  The encoding for optional values uses a 32-bit integer to hold the size of variable-length types. Calling `startSize` writes a placeholder value for the size; after writing the data, call `endSize` to patch the placeholder with the actual size.

- `byte[] finished()`
  Indicates that marshaling is complete and returns the encoded byte sequence. This member function must only be called once.

- `void reset(boolean clearBuffer)`
  Resets the writing position of the stream to the beginning. The boolean argument `clearBuffer` determines whether the stream releases the internal buffer it allocated to hold the encoded data. If `clearBuffer` is true, the stream releases the buffer in order to make it eligible for garbage collection. If `clearBuffer` is false, the stream retains the buffer to avoid generating unnecessary garbage.

- `void writeSerializable(java.io.Serializable v)`
  Writes a serializable Java object to the stream.

- `void destroy()`
  Applications must call this function in order to reclaim resources.

Here is a simple example that demonstrates how to insert a boolean and a sequence of strings into a stream:

Copyright 2016, ZeroC, Inc.

Java

```
final String[] seq = { "Ice", "rocks!" };
Ice.OutputStream out = Ice.Util.createOutputStream(communicator);
try {
    out.writeBool(true);
    out.writeStringSeq(seq);
    byte[] data = out.finished();
} finally {
    out.destroy();
}
```

See Also

- Basic Data Encoding
- Data Encoding for Classes
- Data Encoding for Exceptions
- Serializable Objects in Java

Copyright 2016, ZeroC, Inc.

Intercepting Object Insertion and Extraction in Java

In some situations it may be necessary for an application to intercept the insertion and extraction of Ice objects. For example, the Ice extension for PHP is implemented using Ice for C++ but represents Ice objects as native PHP objects. The PHP extension accomplishes this by manually encoding and decoding Ice objects as directed by the data encoding rules. However, the extension obviously cannot pass a native PHP object to the C++ stream function for writing objects. To bridge this gap between object systems, Ice supplies the helper classes described below.

On this page:

- Inserting Objects in Java
- Extracting Objects in Java

Inserting Objects in Java

The `ObjectWriter` base class facilitates the insertion of objects to a stream:

Java

```
package Ice;

public abstract class ObjectWriter extends ObjectImpl {
    public abstract void write(OutputStream out);
    // ...
}
```

A foreign Ice object is inserted into a stream using the following technique:

1. A Java "wrapper" class is derived from `ObjectWriter`. This class wraps the foreign object and implements the `write` member function.
2. An instance of the wrapper class is passed to `writeObject`. (This is possible because `ObjectWriter` derives from `Ice.Object`.) Eventually, the `write` member function is invoked on the wrapper instance.
3. The implementation of `write` encodes the object's state as directed by the data encoding rules for classes.

It is the application's responsibility to ensure that there is a one-to-one mapping between foreign Ice objects and wrapper objects. This is necessary in order to ensure the proper encoding of object graphs.

Extracting Objects in Java

The `ObjectReader` class facilitates the extraction of objects from a stream:

Copyright 2016, ZeroC, Inc.

Java

```
package Ice;

public abstract class ObjectReader extends ObjectImpl {
    public abstract void read(InputStream in);
    // ...
}

public interface CompactIdResolver {
    String resolve(int id);
}
```

Extracting the state of a foreign Ice object is more complicated than insertion:

1. A Java "wrapper" class is derived from `ObjectReader`. An instance of this class represents a foreign Ice object.
2. An object factory is installed that returns instances of the wrapper class. Note that a single object factory can be used for all Slice types if it is registered with an empty Slice type ID.
3. A Java callback class implements the `ReadObjectCallback` interface. The implementation of `invoke` expects its argument to be either null or an instance of the wrapper class as returned by the object factory.
4. An instance of the callback class is passed to `readObject`.
5. When the stream is ready to extract the state of an object, it invokes `read` on the wrapper class. The implementation of `read` decodes the object's state as directed by the data encoding rules for classes.
6. The callback object passed to `readObject` is invoked, passing the instance of the wrapper object. All other callback objects representing the same instance in the stream (in case of object graphs) are invoked with the same wrapper object.

If your class definitions use compact type IDs, you must also supply an implementation of `CompactIdResolver` when initializing the communicator. This object is responsible for translating numeric type IDs into their string equivalents.

See Also

- Client-Side Slice-to-PHP Mapping
- Data Encoding
- Data Encoding for Classes
- Class Factories in Java
- Communicator Initialization

Copyright 2016, ZeroC, Inc.

Intercepting User Exception Insertion and Extraction in Java

Inserting a User Exception in Java

As in the case of Ice objects, a Dynamic Ice application may represent user exceptions in a native format that is not directly compatible with the Ice API. If the application needs to raise such a user exception to the Ice run time, the exception must be wrapped in a subclass of `Ice.UserException`. The Dynamic Ice API provides a class to simplify this process:

Java

```
package Ice;

public abstract class UserExceptionWriter extends UserException {

    public UserExceptionWriter(Communicator communicator);

    public abstract void write(Ice.OutputStream os);

    // ...
}
```

A subclass of `UserExceptionWriter` is responsible for supplying a communicator to the constructor, and for implementing the following methods:

- `void write(OutputStream os)`
  This method is invoked when the Ice run time is ready to marshal the exception. The subclass must marshal the exception using the encoding rules for exceptions.

Extracting a User Exception in Java

An application extracts a user exception by calling one of two versions of the `throwException` method defined in the `InputStream` class:

Java

```
package Ice;

public interface InputStream {
    void throwException() throws UserException;
    void throwException(UserExceptionReaderFactory factory) throws UserException;

    // ...
}
```

The version without any arguments attempts to locate and throw a Java implementation of the encoded exception using classes generated by the Slice-to-Java compiler.

If your goal is to create an exception in another type system, such as a native PHP exception object, you must call the second version of `throwException` and pass an implementation of `UserExceptionReaderFactory`:

Copyright 2016, ZeroC, Inc.

Java

```
package Ice;

public interface UserExceptionReaderFactory {
    void createAndThrow(String typeId) throws UserException;
}
```

As the stream iterates over slices of an exception from most-derived to least-derived, it invokes `createAndThrow` passing the type ID of each slice, giving the application an opportunity to raise an instance of `UserExceptionReader`:

Java

```
package Ice;

public abstract class UserExceptionReader extends UserException {
    protected UserExceptionReader(Communicator communicator);

    public abstract void read(InputStream is);

    public abstract String ice_name();

    protected Communicator _communicator;
}
```

Subclasses of `UserExceptionReader` must implement the abstract functions. In particular, the implementation of `read` must call `InputStream.startException`, unmarshal the remaining slices, and then call `InputStream.endException`.

See Also

- Intercepting Object Insertion and Extraction in Java
- Data Encoding for Exceptions

Copyright 2016, ZeroC, Inc.

Stream Helper Methods in Java

The stream classes provide all of the low-level methods necessary for encoding and decoding Ice types. However, it would be tedious and error-prone to manually encode complex Ice types such as classes, structs, and dictionaries using these low-level functions. For this reason, the Slice compiler optionally generates helper methods for streaming complex Ice types.

We will use the following Slice definitions to demonstrate the language mapping:

Slice

```
module M {
    sequence<...> Seq;
    dictionary<...> Dict;
    struct S {
        ...
    };
    enum E { ... };
    class C {
        ...
    };
};
```

The Slice compiler generates the corresponding helper methods shown below:

Copyright 2016, ZeroC, Inc.

Java

```
package M;

public class SeqHelper {
    public static T[] read(Ice.InputStream in);
    public static void write(Ice.OutputStream out, T[] v);
    public static void Ice.OptionalFormat optionalFormat();
}

public class DictHelper {
    public static java.util.Map read(Ice.InputStream in);
    public static void write(Ice.OutputStream out, java.util.Map<..., ...> v);
    public static void Ice.OptionalFormat optionalFormat();
}

public class SHelper {
    public static S read(Ice.InputStream in);
    public static void write(Ice.OutputStream out, S v);
    public static void Ice.OptionalFormat optionalFormat();
}

public class EHelper {
    public static E read(Ice.InputStream in);
    public static void write(Ice.OutputStream out, E v);
    public static void Ice.OptionalFormat optionalFormat();
}

public class CHelper {
    public static void read(Ice.InputStream in, CHolder h);
    public static void write(Ice.OutputStream out, C v);
    public static void Ice.OptionalFormat optionalFormat();
}

public class CPrxHelper {
    public static CPrx read(Ice.InputStream in);
    public static void write(Ice.OutputStream out, CPrx v);
    public static void Ice.OptionalFormat optionalFormat();
}
```

The `optionalFormat` method simplifies the task of writing optional values.

The Slice compiler also generates the following methods for `struct` and `enum` types:

Copyright 2016, ZeroC, Inc.

Java

```
public class S ... {
    ...
    public void ice_read(Ice.InputStream in);
    public void ice_write(Ice.OutputStream out);
};
public class E... {
    ...
    public void ice_read(Ice.InputStream in);
    public void ice_write(Ice.OutputStream out);
}
```

Be aware that a call to `CHelper.read` does not result in the immediate extraction of an Ice object. The `value` member of the given `CHolder` object is updated when `readPendingObjects` is invoked on the input stream.

See Also

- Data Encoding
- slice2java Command-Line Options
- The InputStream Interface in Java
- Optional Values

Copyright 2016, ZeroC, Inc.

C-Sharp Streaming Interfaces

We discuss the stream classes first, followed by the helper functions, and finish with an advanced use case.

Topics

- The InputStream Interface in C-Sharp
- The OutputStream Interface in C-Sharp
- Stream Helper Methods in C-Sharp
- Intercepting Object Insertion and Extraction in C-Sharp
- Intercepting User Exception Insertion and Extraction in C-Sharp

Copyright 2016, ZeroC, Inc.

The InputStream Interface in C-Sharp

An `InputStream` is created using the following function:

C#

```
namespace Ice
{
    public sealed class Util
    {
        public static InputStream createInputStream(
            Communicator communicator,
            byte[] bytes);

        public static InputStream createInputStream(
            Communicator communicator,
            byte[] bytes,
            EncodingVersion version);

        public static InputStream wrapInputStream(
            Communicator communicator,
            byte[] bytes);

        public static InputStream wrapInputStream(
            Communicator communicator,
            byte[] bytes,
            EncodingVersion version);
    }
}
```

You can optionally specify an encoding version for the stream, otherwise the stream uses the communicator's default encoding version.

Note that the `createInputStream` functions make a copy of the supplied data, whereas the `wrapInputStream` functions avoid copies by keeping a reference to the data. If you use `wrapInputStream`, it is your responsibility to ensure that the memory buffer remains unmodified for the lifetime of the `InputStream` object.

The `InputStream` interface is shown below.

C#

```
namespace Ice
{
    public interface InputStream
    {
        Communicator communicator();

        void sliceObjects(bool slice);

        bool readBool();
        bool[] readBoolSeq();

        byte readByte();
        byte[] readByteSeq();
```


Copyright 2016, ZeroC, Inc.

```
short readShort();
short[] readShortSeq();

int readInt();
int[] readIntSeq();

long readLong();
long[] readLongSeq();

float readFloat();
float[] readFloatSeq();

double readDouble();
double[] readDoubleSeq();

string readString();
string[] readStringSeq();

int readSize();
int readAndCheckSeqSize(int minSize);

ObjectPrx readProxy();

void readObject(ReadObjectCallback cb);

int readEnum(int maxValue);

void throwException();
void throwException(UserExceptionReaderFactory factory);

void startObject();
SlicedData endObject(bool preserve);

void startException();
SlicedData endException(bool preserve);

string startSlice();
void endSlice();
void skipSlice();

EncodingVersion startEncapsulation();
void endEncapsulation();
EncodingVersion skipEncapsulation();
EncodingVersion getEncoding();

void readPendingObjects();

object readSerializable();

void rewind();

void skip(int sz);
void skipSize();

bool readOptional(int tag, OptionalFormat format);
```


Copyright 2016, ZeroC, Inc.

```
        int pos();

        void destroy();
```


Copyright 2016, ZeroC, Inc.

```
    }
}
```

Member functions are provided for extracting all of the primitive types, as well as sequences of primitive types; these are self-explanatory. The remaining member functions have the following semantics:

- `void sliceObjects(boolean slice)`
  Determines the behavior of the stream when extracting Ice objects. An Ice object is "sliced" when a factory cannot be found for a Slice type ID, resulting in the creation of an object of a less-derived type. Slicing is typically disabled when the application expects all object factories to be present, in which case the exception `NoObjectFactoryException` is raised. The default behavior is to allow slicing.

- `int readSize()`
  The Ice encoding has a compact representation to indicate size. This function extracts a size and returns it as an integer.

- `int readAndCheckSeqSize(int minWireSize)`
  Like `readSize`, this function reads a size and returns it, but also verifies that there is enough data remaining in the unmarshaling buffer to successfully unmarshal the elements of the sequence. The `minWireSize` parameter indicates the smallest possible on-the-wire representation of a single sequence element. If the unmarshaling buffer contains insufficient data to unmarshal the sequence, the function throws `UnmarshalOutOfBoundsException`.

- `Ice.ObjectPrx readProxy()`
  This function returns an instance of the base proxy type, `ObjectPrx`. The Slice compiler optionally generates helper functions to extract proxies of user-defined types.

- `void readObject(ReadObjectCallback cb)`
  The Ice encoding for class instances requires extraction to occur in stages. The `readObject` function accepts a callback object of type `ReadObjectCallback`, whose definition is shown below:

  C#

  ```
  namespace Ice {
      public interface ReadObjectCallback     {
          void invoke(Ice.Object obj);
      }
  }
  ```

  When the object instance is available, the callback object's `invoke` member function is called. The application must call `readPendingObjects` to ensure that all instances are properly extracted. Note that applications rarely need to invoke this member function directly; the helper functions generated by the Slice compiler are easier to use.

- `int readEnum(int maxValue)`
  Unmarshals the integer value of an enumerator. The `maxValue` argument represents the highest enumerator value in the enumeration. Consider the following definitions:

  Slice

  ```
  enum Color { red, green, blue };
  enum Fruit { Apple, Pear=3, Orange };
  ```

  The maximum value for `Color` is 2, and the maximum value for `Fruit` is 4.

- `void throwException()`
  `void throwException(UserExceptionReaderFactory factory)`
  These functions extract a user exception from the stream and throw it. If the stored exception is of an unknown type, the functions

Copyright 2016, ZeroC, Inc.

attempt to extract and throw a less-derived exception. If that also fails, an exception is thrown: for the 1.0 encoding, the exception is `UnmarshalOutOfBoundsException`, for the 1.1 encoding, the exception is `UnknownUserException`.

- `void startObject()`
  `SlicedData endObject(bool preserve)`
  The `startObject` method must be called prior to reading the slices of an object. The `endObject` method must be called after all slices have been read. Pass true to `endObject` in order to preserve the slices of any unknown more-derived types, or false to discard the slices. If `preserve` is true and the stream actually preserved any slices, the return value of `endObject` is a non-nil `SlicedData` object that encapsulates the slice data. If the caller later wishes to forward the object with any preserved slices intact, it must supply this `SlicedData` object to the output stream.

- `void startException()`
  `SlicedData endException(bool preserve)`
  The `startException` method must be called prior to reading the slices of an exception. The `endException` method must be called after all slices have been read. Pass true to `endException` in order to preserve the slices of any unknown more-derived types, or false to discard the slices. If `preserve` is true and the stream actually preserved any slices, the return value of `endException` is a non-nil `SlicedData` object that encapsulates the slice data. If the caller later wishes to forward the exception with any preserved slices intact, it must supply this `SlicedData` object to the output stream.

- `string startSlice()`
  `void endSlice()`
  `void skipSlice()`
  Start, end, and skip a slice of member data, respectively. These functions are used when manually extracting the slices of an Ice object or user exception. The `startSlice` method returns the type ID of the next slice, which may be an empty string depending on the format used to encode the object or exception.

- `EncodingVersion startEncapsulation()`
  `void endEncapsulation()`
  `EncodingVersion skipEncapsulation()`
  Start, end, and skip an encapsulation, respectively. The `startEncapsulation` and `skipEncapsulation` methods return the encoding version used to encode the contents of the encapsulation.

- `EncodingVersion getEncoding()`
  Returns the encoding version currently in use by the stream.

- `void readPendingObjects()`
  An application must call this function after all other data has been extracted, but only if Ice objects were encoded. This function extracts the state of Ice objects and invokes their corresponding callback objects (see `readObject`). Note that `endEncapsulation` implicitly calls `readPendingObjects` if necessary.

- `void readPendingObjects()`
  An application must call this function after all other data has been extracted, but only if Ice objects were encoded. This function extracts the state of Ice objects and invokes their corresponding callback objects (see `readObject`). For backward compatibility with encoding version 1.0, this function must only be called when non-optional data members or parameters use class types.

- `object readSerializable()`
  Reads a serializable .NET object from the stream.

- `void rewind()`
  Resets the position of the stream to the beginning.

- `void skip(int sz)`
  Skips the given number of bytes.

- `void skipSize()`
  Reads a size at the current position and skips that number of bytes.

- `bool readOptional(int tag, OptionalFormat fmt)`
  Returns true if an optional value with the given tag and format is present, or false otherwise. If this method returns true, the data associated with that optional value must be read next. Optional values must be read in order by tag from least to greatest. The `OptionalFormat` enumeration is defined as follows:

Copyright 2016, ZeroC, Inc.

C#

```
namespace Ice {
    enum OptionalFormat {
        OptionalFormatF1, OptionalFormatF2, OptionalFormatF4,
OptionalFormatF8,
        OptionalFormatSize, OptionalFormatVSize, OptionalFormatFSize,
        OptionalFormatEndMarker
    }
}
```

Refer to the encoding discussion for more information on the meaning of these values.

- `void destroy()`
  Applications must call this function in order to reclaim resources.

Here is a simple example that demonstrates how to extract a boolean and a sequence of strings from a stream:

C#

```
byte[] data = ...
Ice.InputStream inStream =
    Ice.Util.createInputStream(communicator, data);
try {
    bool b = inStream.readBool();
    string[] seq = inStream.readStringSeq();
} finally {
    inStream.destroy();
}
```

See Also

- Basic Data Encoding
- Data Encoding for Classes
- Data Encoding for Exceptions
- Serializable Objects in C-Sharp
- Stream Helper Methods in C-Sharp

Copyright 2016, ZeroC, Inc.

The OutputStream Interface in C-Sharp

An `OutputStream` is created using the following function:

C#

```
namespace Ice
{
    public sealed class Util
    {
        public static OutputStream
        createOutputStream(Communicator communicator);

        public static OutputStream
        createOutputStream(Communicator communicator, EncodingVersion version);
    }
}
```

You can optionally specify an encoding version for the stream, otherwise the stream uses the communicator's default encoding version.

The `OutputStream` class is shown below.

C#

```
namespace Ice
{
    public interface OutputStream
    {
        Communicator communicator();

        void writeBool(bool v);
        void writeBoolSeq(bool[] v);

        void writeByte(byte v);
        void writeByteSeq(byte[] v);

        void writeShort(short v);
        void writeShortSeq(short[] v);

        void writeInt(int v);
        void writeIntSeq(int[] v);

        void writeLong(long v);
        void writeLongSeq(long[] v);

        void writeFloat(float v);
        void writeFloatSeq(float[] v);

        void writeDouble(double v);
        void writeDoubleSeq(double[] v);

        void writeString(string v);
```


Copyright 2016, ZeroC, Inc.

```
    void writeStringSeq(string[] v);

    void writeSize(int sz);

    void writeProxy(ObjectPrx v);

    void writeObject(Ice.Object v);

    void writeEnum(int v, int maxValue);

    void writeException(UserException ex);

    void startObject(SlicedData sd);
    void endObject();

    void startException(SlicedData sd);
    void endException();

    void startSlice(string typeId, bool last);
    void endSlice();

    void startEncapsulation(EncodingVersion encoding, FormatType format);
    void startEncapsulation();
    void endEncapsulation();

    EncodingVersion getEncoding();

    void writePendingObjects();

    bool writeOptional(int tag, OptionalFormat format);

    int pos();

    void rewrite(int sz, int pos);

    void startSize();
    void endSize();

    byte[] finished();

    void reset(bool clearBuffer);

    void writeSerializable(object v);

    void destroy();
```


Copyright 2016, ZeroC, Inc.

```
    }
}
```

Member functions are provided for inserting all of the primitive types, as well as sequences of primitive types; these are self-explanatory. The remaining member functions have the following semantics:

- `void writeSize(int sz)`
  The Ice encoding has a compact representation to indicate size. This function converts the given non-negative integer into the proper encoded representation.
- `void writeObject(Ice.Object v)`
  Inserts an Ice object. The Ice encoding for class instances may cause the insertion of this object to be delayed, in which case the stream retains a reference to the given object and does not insert its state it until `writePendingObjects` is invoked on the stream.
- `void writeProxy(Ice.ObjectPrx v)`
  Inserts a proxy.
- `void writeEnum(int val, int maxValue)`
  Writes the integer value of an enumerator. The `maxValue` argument represents the highest enumerator value in the enumeration. Consider the following definitions:

  Slice

  ```
  enum Color { red, green, blue };
  enum Fruit { Apple, Pear=3, Orange };
  ```

  The maximum value for `Color` is 2, and the maximum value for `Fruit` is 4.
- `void writeException(UserException ex)`
  Inserts a user exception. The exception argument may be an instance of `UserExceptionWriter`, which allows you to implement your own exception marshaling logic.
- `void startObject(SlicedData sd)`
  `void endObject()`
  When marshaling the slices of an object, the application must first call `startObject`, then marshal the slices, and finally call `endObject`. The caller can pass a `SlicedData` object containing the preserved slices of unknown more-derived types, or 0 if there are no preserved slices.
- `void startException(SlicedData sd)`
  `void endException()`
  When marshaling the slices of an exception, the application must first call `startException`, then marshal the slices, and finally call `endException`. The caller can pass a `SlicedData` object containing the preserved slices of unknown more-derived types, or 0 if there are no preserved slices.
- `void startSlice(string typeId, bool last)`
  `void endSlice()`
  Starts and ends a slice of object or exception member data. The call to `startSlice` must include the type ID for the current slice, and a boolean indicating whether this is the last slice of the object or exception.
- `void startEncapsulation(EncodingVersion encoding, FormatType format)`
  `void startEncapsulation()`
  `void endEncapsulation()`
  Starts and ends an encapsulation, respectively. The first overloading of `startEncapsulation` allows you to specify the encoding version as well as the format to use for any objects and exceptions marshaled within this encapsulation.
- `EncodingVersion getEncoding()`
  Returns the encoding version currently being used by the stream.
- `void writePendingObjects()`
  Encodes the state of Ice objects whose insertion was delayed during `writeObject`. This member function must only be called once. For backward compatibility with encoding version 1.0, this function must only be called when non-optional data members or

Copyright 2016, ZeroC, Inc.

parameters use class types.

- `bool writeOptional(int tag, OptionalFormat fmt)`
  Prepares the stream to write an optional value with the given tag and format. Returns true if the value should be written, or false otherwise. A return value of false indicates that the encoding version in use by the stream does not support optional values. If this method returns true, the data associated with that optional value must be written in order by tag from least to greatest. The `OptionalFormat` enumeration is defined as follows:

  C#

  ```
  namespace Ice {
      enum OptionalFormat {
          OptionalFormatF1, OptionalFormatF2, OptionalFormatF4,
  OptionalFormatF8,
          OptionalFormatSize, OptionalFormatVSize, OptionalFormatFSize,
          OptionalFormatEndMarker
      }
  }
  ```

  Refer to the encoding discussion for more information on the meaning of these values.

- `int pos()`
  `void rewrite(int sz, int pos)`
  The `pos` method returns the stream's current position, and `rewrite` allows you to overwrite a 32-bit integer value at the given position in the stream. Calling `rewrite` does not change the stream's current position.

- `byte[] finished()`
  Indicates that marshaling is complete and returns the encoded byte sequence. This member function must only be called once.

- `void reset(bool clearBuffer)`
  Resets the writing position of the stream to the beginning. The boolean argument `clearBuffer` determines whether the stream releases the internal buffer it allocated to hold the encoded data. If `clearBuffer` is true, the stream releases the buffer in order to make it eligible for garbage collection. If `clearBuffer` is false, the stream retains the buffer to avoid generating unnecessary garbage.

- `void writeSerializable(object v)`
  Writes a serializable .NET object to the stream.

- `void destroy()`
  Applications must call this function in order to reclaim resources.

Here is a simple example that demonstrates how to insert a boolean and a sequence of strings into a stream:

C#

```
string[] seq = { "Ice", "rocks!" };
Ice.OutputStream outStream
    = Ice.Util.createOutputStream(communicator);
try {
    outStream.writeBool(true);
    outStream.writeStringSeq(seq);
    byte[] data = outStream.finished();
} finally {
    outStream.destroy();
}
```


Copyright 2016, ZeroC, Inc.

See Also

- Basic Data Encoding
- Data Encoding for Classes
- Data Encoding for Exceptions
- Serializable Objects in C-Sharp

Copyright 2016, ZeroC, Inc.

Stream Helper Methods in C-Sharp

The stream classes provide all of the low-level methods necessary for encoding and decoding Ice types. However, it would be tedious and error-prone to manually encode complex Ice types such as classes, structs, and dictionaries using these low-level functions. For this reason, the Slice compiler optionally generates helper methods for streaming complex Ice types.

We will use the following Slice definitions to demonstrate the language mapping:

Slice

```
module M {
    sequence<...> Seq;
    dictionary<...> Dict;
    struct S {
        ...
    };
    enum E { ... };
    class C {
        ...
    };
};
```

The Slice compiler generates the corresponding helper methods shown below:

Copyright 2016, ZeroC, Inc.

C#

```
namespace M
{
    public sealed class SeqHelper
    {
        public static int[] read(Ice.InputStream _in);
        public static void write(Ice.OutputStream _out, int[] _v);
        public static Ice.OptionalFormat optionalFormat();
    }

    public sealed class DictHelper
    {
        public static Dictionary<...> read(Ice.InputStream _in);
        public static void write(Ice.OutputStream _out, Dictionary<...> _v);
        public static Ice.OptionalFormat optionalFormat();
    }

    public sealed class SHelper
    {
        public static S read(Ice.InputStream _in);
        public static void write(Ice.OutputStream _out, S _v);
        public static Ice.OptionalFormat optionalFormat();
    }

    public sealed class EHelper
    {
        public static M.E read(Ice.InputStream _in);
        public static void write(Ice.OutputStream _out, M.E _v);
        public static Ice.OptionalFormat optionalFormat();
    }

    public sealed class CHelper
    {
        public CHelper(Ice.InputStream _in);
        public void read();
        public static void write(Ice.OutputStream _out, C _v);
        public M.C value
        {
            get;
        }
        public static Ice.OptionalFormat optionalFormat();
        // ...
    }

    public sealed class CPrxHelper : Ice.ObjectPrxHelperBase, CPrx
    {
        public static CPrx read(Ice.InputStream _in);
        public static void write(Ice.OutputStream _out, CPrx _v);
        public static Ice.OptionalFormat optionalFormat();
    }
}
```


Copyright 2016, ZeroC, Inc.

The `optionalFormat` method simplifies the task of writing optional values.

The Slice compiler also generates the following methods for `struct` types:

C#

```
public struct S {
    ...
    public void ice_read(Ice.InputStream in);
    public void ice_write(Ice.OutputStream out);
}
```

Be aware that a call to `CHelper.read` does not result in the immediate extraction of an Ice object. The `value` property of the given `CHelper` object is updated when `readPendingObjects` is invoked on the input stream.

See Also

- slice2cs Command-Line Options
- The InputStream Interface in C-Sharp
- Optional Values

Copyright 2016, ZeroC, Inc.

Intercepting Object Insertion and Extraction in C-Sharp

In some situations it may be necessary for an application to intercept the insertion and extraction of Ice objects. For example, the Ice extension for PHP is implemented using Ice for C++ but represents Ice objects as native PHP objects. The PHP extension accomplishes this by manually encoding and decoding Ice objects as directed by the data encoding rules. However, the extension obviously cannot pass a native PHP object to the C++ stream function `writeObject`. To bridge this gap between object systems, Ice supplies the helper classes described below.

On this page:

- Inserting Objects in C#
- Extracting Objects in C#

Inserting Objects in C#

The `ObjectWriter` base class facilitates the insertion of objects to a stream:

C#

```
namespace Ice
{
    public abstract class ObjectWriter : ObjectImpl
    {
        public abstract void write(OutputStream outStream);
        // ...
    }
}
```

A foreign Ice object is inserted into a stream using the following technique:

1. A C# "wrapper" class is derived from `ObjectWriter`. This class wraps the foreign object and implements the `write` member function.
2. An instance of the wrapper class is passed to `writeObject`. (This is possible because `ObjectWriter` derives from `Ice.Object`.) Eventually, the `write` member function is invoked on the wrapper instance.
3. The implementation of `write` encodes the object's state as directed by the data encoding rules for classes.

It is the application's responsibility to ensure that there is a one-to-one mapping between foreign Ice objects and wrapper objects. This is necessary in order to ensure the proper encoding of object graphs.

Extracting Objects in C#

The `ObjectReader` class facilitates the extraction of objects from a stream:

Copyright 2016, ZeroC, Inc.

C#

```
namespace Ice
{
    public abstract class ObjectReader : ObjectImpl
    {
        public abstract void read(InputStream inStream);
        // ...
    }

    public delegate string CompactIdResolver(int id);
}
```

Extracting the state of a foreign Ice object is more complicated than insertion:

1. A C# "wrapper" class is derived from `ObjectReader`. An instance of this class represents a foreign Ice object.
2. An object factory is installed that returns instances of the wrapper class. Note that a single object factory can be used for all Slice types if it is registered with an empty Slice type ID.
3. A C# callback class implements the `ReadObjectCallback` interface. The implementation of `invoke` expects its argument to be either null or an instance of the wrapper class as returned by the object factory.
4. An instance of the callback class is passed to `readObject`.
5. When the stream is ready to extract the state of an object, it invokes `read` on the wrapper class. The implementation of `read` decodes the object's state as directed by the data encoding rules for classes.
6. The callback object passed to `readObject` is invoked, passing the instance of the wrapper object. All other callback objects representing the same instance in the stream (in case of object graphs) are invoked with the same wrapper object.

If your class definitions use compact type IDs, you must also supply an implementation of `CompactIdResolver` when initializing the communicator. This object is responsible for translating numeric type IDs into their string equivalents.

See Also

- Client-Side Slice-to-PHP Mapping
- Data Encoding
- Data Encoding for Classes
- Class Factories in C#
- Communicator Initialization

Copyright 2016, ZeroC, Inc.

Intercepting User Exception Insertion and Extraction in C-Sharp

Inserting a User Exception in C#

As in the case of Ice objects, a Dynamic Ice application may represent user exceptions in a native format that is not directly compatible with the Ice API. If the application needs to raise such a user exception to the Ice run time, the exception must be wrapped in a subclass of `Ice.UserException`. The Dynamic Ice API provides a class to simplify this process:

C#

```
namespace Ice
{
    public abstract class UserExceptionWriter : UserException
    {
        public UserExceptionWriter(Communicator communicator);

        public abstract void write(OutputStream os);

        // ...
    }
}
```

A subclass of `UserExceptionWriter` is responsible for supplying a communicator to the constructor, and for implementing the following methods:

- `void write(OutputStream os)`
  This method is invoked when the Ice run time is ready to marshal the exception. The subclass must marshal the exception using the encoding rules for exceptions.

Extracting a User Exception in C#

An application extracts a user exception by calling one of two versions of the `throwException` method defined in the `InputStream` class:

C#

```
namespace Ice {
    public interface InputStream {
        void throwException();
        void throwException(UserExceptionReaderFactory factory);

        // ...
    }
}
```

The version without any arguments attempts to locate and throw a C# implementation of the encoded exception using classes generated by the Slice-to-C# compiler.

If your goal is to create an exception in another type system, such as a native PHP exception object, you must call the second version of `throwException` and pass an implementation of `UserExceptionReaderFactory`:

Copyright 2016, ZeroC, Inc.

C#

```
namespace Ice {
    public interface UserExceptionReaderFactory {
        void createAndThrow(string typeId);
    }
}
```

As the stream iterates over slices of an exception from most-derived to least-derived, it invokes `createAndThrow` passing the type ID of each slice, giving the application an opportunity to raise an instance of `UserExceptionReader`:

C#

```
namespace Ice {
    public abstract class UserExceptionReader : UserException {
        protected UserExceptionReader(Communicator communicator);

        public abstract void read(InputStream is);

        public abstract string ice_name();

        protected Communicator _communicator;
    }
}
```

Subclasses of `UserExceptionReader` must implement the abstract functions. In particular, the implementation of `read` must call `InputStream.startException`, unmarshal the remaining slices, and then call `InputStream.endException`.

See Also

- Intercepting Object Insertion and Extraction in C-Sharp
- Data Encoding for Exceptions

Copyright 2016, ZeroC, Inc.

Dynamic Invocation and Dispatch

Topics

- Dynamic Invocation and Dispatch Overview
- Dynamic Invocation and Dispatch in C++
- Dynamic Invocation and Dispatch in Java
- Dynamic Invocation and Dispatch in C-Sharp

Copyright 2016, ZeroC, Inc.

Dynamic Invocation and Dispatch Overview

On this page:

- Use Cases for Dynamic Invocation and Dispatch
- Dynamic Invocation using ice_invoke
- Dynamic Dispatch using Blobject

Use Cases for Dynamic Invocation and Dispatch

Ice applications generally use the static invocation model, in which the application invokes a Slice operation by calling a member function on a generated proxy class. In the server, the static dispatch model behaves similarly: the request is dispatched to the servant as a statically-typed call to a member function. Underneath this statically-typed facade, the Ice run times in the client and server are exchanging sequences of bytes representing the encoded request arguments and results. These interactions are illustrated below:

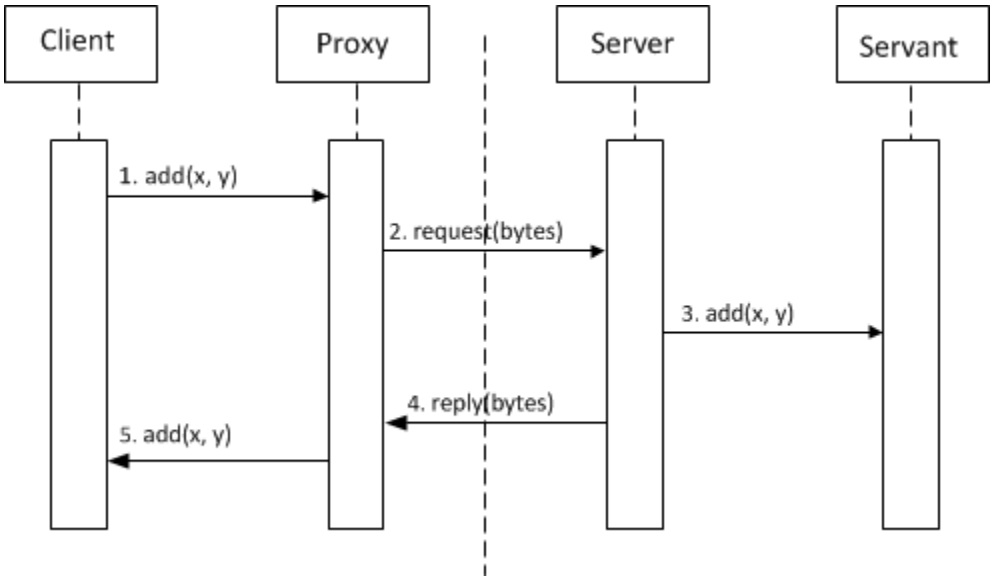

Interactions in a static invocation.

1. The client initiates a call to the Slice operation `add` by calling the member function `add` on a proxy.
2. The generated proxy class marshals the arguments into a sequence of bytes and transmits them to the server.
3. In the server, the generated servant class unmarshals the arguments and calls `add` on the subclass.
4. The servant marshals the results and returns them to the client.
5. Finally, the client's proxy unmarshals the results and returns them to the caller.

The application is blissfully unaware of this low-level machinery, and in the majority of cases that is a distinct advantage. In some situations, however, an application can leverage this machinery to accomplish tasks that are not possible in a statically-typed environment. Ice provides the dynamic invocation and dispatch models for these situations, allowing applications to send and receive requests as encoded sequences of bytes instead of statically-typed arguments.

The dynamic invocation and dispatch models offer several unique advantages to Ice services that forward requests from senders to receivers, such as Glacier2 and IceStorm. For these services, the request arguments are an opaque byte sequence that can be forwarded without the need to unmarshal and remarshal the arguments. Not only is this significantly more efficient than a statically-typed implementation, it also allows intermediaries such as Glacier2 and IceStorm to be ignorant of the Slice types in use by senders and receivers.

Another use case for the dynamic invocation and dispatch models is scripting language integration. The Ice extensions for Python, PHP, and Ruby invoke Slice operations using the dynamic invocation model; the request arguments are encoded using the streaming interfaces.

It may be difficult to resist the temptation of using a feature like dynamic invocation or dispatch, but we recommend that you carefully consider the risks and complexities of such a decision. For example, an application that uses the streaming interface to manually encode and decode request arguments has a high risk of failure if the argument signature of an operation changes. In contrast, this risk is greatly reduced in the static invocation and dispatch models because errors in a strongly-typed language are found early, during compilation. Therefore, we caution you against using this capability except where its advantages significantly outweigh the risks.

Dynamic Invocation using `ice_invoke`

Copyright 2016, ZeroC, Inc.

Dynamic invocation is performed using the proxy member function `ice_invoke`, defined in the proxy base class `ObjectPrx`. If we were to define the function in Slice, it would look like this:

Slice

```
sequence<byte> ByteSeq;

bool ice_invoke(
    string operation,
    Ice::OperationMode mode,
    ByteSeq inParams,
    out ByteSeq outParams
);
```

The first argument is the name of the Slice operation.

This is the Slice name of the operation, not the name as it might be mapped to any particular language. For example, the string `"while"` is the name of the Slice operation `while`, and not `"_cpp_while"` (C++) or `"_while"` (Java).

The second argument is an enumerator from the Slice type `Ice::OperationMode`; the possible values are `Normal` and `Idempotent`. The third argument, `inParams`, represents an encapsulation of the encoded in-parameters of the operation.

A return value of `true` indicates a successful invocation, in which case an encapsulation of the marshaled form of the operation's results (if any) is provided in `outParams`. A return value of `false` signals the occurrence of a user exception whose encapsulated data is provided in `outParams`. The caller must also be prepared to catch local exceptions, which are thrown directly.

Dynamic Dispatch using `Blobject`

A server enables dynamic dispatch by creating a subclass of `Blobject` (the name is derived from *blob*, meaning a blob of bytes). The Slice equivalent of `Blobject` is shown below:

Slice

```
sequence<byte> ByteSeq;

interface Blobject {
    bool ice_invoke(ByteSeq inParams, out ByteSeq outParams);
};
```

The `inParams` argument supplies an encapsulation of the encoded in-parameters. The contents of the `outParams` argument depends on the outcome of the invocation: if the operation succeeded, `ice_invoke` must return `true` and place an encapsulation of the encoded results in `outParams`; if a user exception occurred, `ice_invoke` must return `false`, in which case `outParams` contains an encapsulation of the encoded exception. The operation may also raise local exceptions such as `OperationNotExistException`.

The language mappings add a trailing argument of type `Ice::Current` to `ice_invoke`, and this provides the implementation with the name of the operation being dispatched.

Because `Blobject` derives from `Object`, an instance is a regular Ice servant just like instances of the classes generated for user-defined Slice interfaces. The primary difference is that all operation invocations on a `Blobject` instance are dispatched through the `ice_invoke` m ember function.

If a `Blobject` subclass intends to decode the in-parameters (and not simply forward the request to another object), then the implementation obviously must know the signatures of all operations it supports. How a `Blobject` subclass determines its type information is an implementation detail that is beyond the scope of this manual.

Note that a `Blobject` servant is also useful if you want to create a message forwarding service, such as Glacier2. In this case, there is no

Copyright 2016, ZeroC, Inc.

need to decode any parameters; instead, the implementation simply forwards each request unchanged to a new destination. You can register a `Blobject` servant as a default servant to easily achieve this.

See Also

- Collocated Invocation and Dispatch
- The Current Object
- Default Servants
- Streaming Interfaces
- Glacier2
- IceStorm

Copyright 2016, ZeroC, Inc.

Dynamic Invocation and Dispatch in C++

This page describes the C++ mapping for the `ice_invoke` proxy function and the `Blobject` class.

On this page:

- ice_invoke in C++
- Using Streams with ice_invoke in C++
- Subclassing Blobject in C++
- Using the Array Mapping for ice_invoke and Blobject in C++

`ice_invoke` in C++

The mapping for `ice_invoke` is shown below:

C++

```
bool ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::vector< Ice::Byte >& inParams,
    std::vector< Ice::Byte >& outParams
);
```

Another overloading of `ice_invoke` (not shown) adds a trailing argument of type `Ice::Context`.

As an example, the code below demonstrates how to invoke the operation `op`, which takes no `in` parameters:

C++

```
Ice::ObjectPrx proxy = ...
try {
    std::vector<Ice::Byte> inParams, outParams;
    if (proxy->ice_invoke("op", Ice::Normal, inParams, outParams)) {
        // Handle success
    } else {
        // Handle user exception
    }
} catch (const Ice::LocalException& ex) {
    // Handle exception
}
```

As a convenience, the Ice run time accepts an empty byte sequence when there are no input parameters and internally translates it into an empty encapsulation. In all other cases, the value for `inParams` must be an encapsulation of the encoded parameters.

Using Streams with `ice_invoke` in C++

The streaming interfaces provide the tools an application needs to dynamically invoke operations with arguments of any Slice type. Consider the following Slice definition:

Copyright 2016, ZeroC, Inc.

Slice

```
module Calc {
    exception Overflow {
        int x;
        int y;
    };
    interface Compute {
        idempotent int add(int x, int y)
            throws Overflow;
    };
};
```

Now let's write a client that dynamically invokes the `add` operation:

C++

```
Ice::ObjectPrx proxy = ...
try {
    std::vector< Ice::Byte > inParams, outParams;

    Ice::OutputStreamPtr out = Ice::createOutputStream(communicator);
    out->startEncapsulation();
    out->writeInt(100); // x
    out->writeInt(-1);  // y
    out->endEncapsulation();
    out->finished(inParams);

    if (proxy->ice_invoke("add", Ice::Idempotent, inParams, outParams)) {
        // Handle success
        Ice::InputStreamPtr in = Ice::createInputStream(communicator, outParams);
        in->startEncapsulation();
        int result = in->readInt();
        in->endEncapsulation();
        assert(result == 99);
    } else {
        // Handle user exception
    }
} catch (const Ice::LocalException& ex) {
    // Handle exception
}
```

You can see here that the input and output parameters are enclosed in encapsulations.

We neglected to handle the case of a user exception in this example, so let's implement that now. We assume that we have compiled our program with the Slice-generated code, therefore we can call `throwException` on the input stream and catch `Overflow` directly:

Copyright 2016, ZeroC, Inc.

C++

```
if (proxy->ice_invoke("add", Ice::Idempotent, inParams, outParams)) {
    // Handle success
    // ...
} else {
    // Handle user exception
    Ice::InputStreamPtr in = Ice::createInputStream(communicator, outParams);
    try {
        in->startEncapsulation();
        in->throwException();
    } catch (const Calc::Overflow& ex) {
        cout << "overflow while adding " << ex.x << " and " << ex.y << endl;
    } catch (const Ice::UserException& ex) {
        // Handle unexpected user exception
    }
}
```

This is obviously a contrived example: if the Slice-generated code is available, why bother using dynamic dispatch? In the absence of Slice-generated code, the caller would need to manually unmarshal the user exception, which is outside the scope of this manual.

As a defensive measure, the code traps `Ice::UserException`. This could be raised if the Slice definition of `add` is modified to include another user exception but this segment of code did not get updated accordingly.

Subclassing `Blobject` in C++

Implementing the dynamic dispatch model requires writing a subclass of `Ice::Blobject`. We continue using the `Compute` interface to demonstrate a `Blobject` implementation:

C++

```
class ComputeI : public Ice::Blobject {
public:
    virtual bool ice_invoke(
        const std::vector<Ice::Byte>& inParams,
        std::vector<Ice::Byte>& outParams,
        const Ice::Current& current);
};
```

An instance of `ComputeI` is an Ice object because `Blobject` derives from `Object`, therefore an instance can be added to an object adapter like any other servant.

For the purposes of this discussion, the implementation of `ice_invoke` handles only the `add` operation and raises `OperationNotExistException` for all other operations. In a real implementation, the servant must also be prepared to receive invocations of the following `Object` operations:

- `string ice_id()`
  Returns the Slice type ID of the servant's most-derived type.
- `StringSeq ice_ids()`
  Returns a sequence of strings representing all of the Slice interfaces supported by the servant, including `"::Ice::Object"`.

Copyright 2016, ZeroC, Inc.

- `bool ice_isA(string id)`
  Returns `true` if the servant supports the interface denoted by the given Slice type ID, or `false` otherwise. This operation is invoked by the proxy function `checkedCast`.
- `void ice_ping()`
  Verifies that the object denoted by the identity and facet contained in `Ice::Current` is reachable.

With that in mind, here is our simplified version of `ice_invoke`:

C++

```
bool ComputeI::ice_invoke(
    const std::vector<Ice::Byte>& inParams,
    std::vector<Ice::Byte>& outParams,
    const Ice::Current& current)
{
    if (current.operation == "add") {
        Ice::CommunicatorPtr communicator = current.adapter->getCommunicator();
        Ice::InputStreamPtr in = Ice::createInputStream(communicator, inParams);
        in->startEncapsulation();
        int x = in->readInt();
        int y = in->readInt();
        in->endEncapsulation();

        Ice::OutputStreamPtr out = Ice::createOutputStream(communicator);
        if (checkOverflow(x, y)) {
            Calc::Overflow ex;
            ex.x = x;
            ex.y = y;
            out->startEncapsulation();
            out->writeException(ex);
            out->endEncapsulation();
            out->finished(outParams);
            return false;
        } else {
            out->startEncapsulation();
            out->writeInt(x + y);
            out->endEncapsulation();
            out->finished(outParams);
            return true;
        }
    } else {
        Ice::OperationNotExistException ex(__FILE__, __LINE__);
        ex.id = current.id;
        ex.facet = current.facet;
        ex.operation = current.operation;
        throw ex;
    }
}
```

If an overflow is detected, the code "raises" the `Calc::Overflow` user exception by calling `writeException` on the output stream and returning `false`, otherwise the return value is encoded and the function returns `true`.

Using the Array Mapping for `ice_invoke` and `Blobject` in C++

Copyright 2016, ZeroC, Inc.

Ice for C++ supports an alternative mapping for `sequence` input parameters that avoids the overhead of extra copying. Since the `ice_invoke` functions treat the encoded input parameters as a value of type `sequence<byte>`, the dynamic invocation and dispatch facility includes interfaces that use the array mapping for the input parameter blob.

Ice provides two overloaded versions of the proxy function `ice_invoke` that use the array mapping. The version that omits the trailing `Ice::Context` argument is shown below:

C++

```
bool ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::pair< const Ice::Byte*, const Ice::Byte* >& in,
    std::vector< Ice::Byte >& out
);
```

A `Blobject` servant uses the array mapping by deriving its implementation class from `Ice::BlobjectArray` and overriding its `ice_invoke` function:

C++

```
class BlobjectArray {
public:
    virtual bool ice_invoke(
        const std::pair<const Ice::Byte*, const Ice::Byte*>& in,
        std::vector<Ice::Byte>& out,
        const Ice::Current& current) = 0;
};
```

See Also

- C++ Mapping for Sequences
- Object Adapters
- Request Contexts
- Streaming Interfaces
- Type IDs
- Object Identity
- Versioning

Copyright 2016, ZeroC, Inc.

Dynamic Invocation and Dispatch in Java

This page describes the Java mapping for the `ice_invoke` proxy function and the `Blobject` class.

On this page:

- ice_invoke in Java
- Using Streams with ice_invoke in Java
- Subclassing Blobject in Java

`ice_invoke` in Java

The mapping for `ice_invoke` is shown below:

Java

```
boolean ice_invoke(
    String operation,
    Ice.OperationMode mode,
    byte[] inParams,
    Ice.ByteSeqHolder outParams
);
```

Another overloading of `ice_invoke` (not shown) adds a trailing argument of type `Ice.Context`.

As an example, the code below demonstrates how to invoke the operation `op`, which takes no `in` parameters:

Java

```
Ice.ObjectPrx proxy = ...
try {
    Ice.ByteSeqHolder outParams = new Ice.ByteSeqHolder();
    if (proxy.ice_invoke("op", Ice.OperationMode.Normal, null, outParams)) {
        // Handle success
    } else {
        // Handle user exception
    }
} catch (Ice.LocalException ex) {
    // Handle exception
}
```

As a convenience, the Ice run time accepts an empty or null byte sequence when there are no input parameters and internally translates it into an empty encapsulation. In all other cases, the value for `inParams` must be an encapsulation of the encoded parameters.

Using Streams with `ice_invoke` in Java

The streaming interfaces provide the tools an application needs to dynamically invoke operations with arguments of any Slice type. Consider the following Slice definition:

Copyright 2016, ZeroC, Inc.

Slice

```
module Calc {
    exception Overflow {
        int x;
        int y;
    };
    interface Compute {
        idempotent int add(int x, int y)
            throws Overflow;
    };
};
```

Now let's write a client that dynamically invokes the `add` operation:

Java

```
Ice.ObjectPrx proxy = ...
try {
    Ice.OutputStream out = Ice.Util.createOutputStream(communicator);
    out.startEncapsulation();
    out.writeInt(100); // x
    out.writeInt(-1);  // y
    out.endEncapsulation();
    byte[] inParams = out.finished();

    Ice.ByteSeqHolder outParams = new Ice.ByteSeqHolder();
    if (proxy.ice_invoke("add", Ice.OperationMode.Idempotent,
inParams, outParams)) {
        // Handle success
        Ice.InputStream in = Ice.Util.createInputStream(communicator,
outParams.value);
        in.startEncapsulation();
        int result = in.readInt();
        in.endEncapsulation();
        assert(result == 99);
    } else {
        // Handle user exception
    }
} catch (Ice.LocalException ex) {
    // Handle exception
}
```

You can see here that the input and output parameters are enclosed in encapsulations.

We neglected to handle the case of a user exception in this example, so let's implement that now. We assume that we have compiled our program with the Slice-generated code, therefore we can call `throwException` on the input stream and catch `Overflow` directly:

Copyright 2016, ZeroC, Inc.

Java

```
    if (proxy.ice_invoke("add", Ice.OperationMode.Idempotent,
inParams, outParams)) {
        // Handle success
        // ...
    } else {
        // Handle user exception
        Ice.InputStream in = Ice.Util.createInputStream(communicator,
outParams.value);
        try {
            in.startEncapsulation();
            in.throwException();
        } catch (Calc.Overflow ex) {
            System.out.println("overflow while adding " + ex.x + " and " + ex.y);
        } catch (Ice.UserException ex) {
            // Handle unexpected user exception
        }
    }
```

This is obviously a contrived example: if the Slice-generated code is available, why bother using dynamic dispatch? In the absence of Slice-generated code, the caller would need to manually unmarshal the user exception, which is outside the scope of this manual.

As a defensive measure, the code traps `Ice.UserException`. This could be raised if the Slice definition of `add` is modified to include another user exception but this segment of code did not get updated accordingly.

Subclassing `Blobject` in Java

Implementing the dynamic dispatch model requires writing a subclass of `Ice.Blobject`. We continue using the `Compute` interface to demonstrate a `Blobject` implementation:

Java

```
public class ComputeI extends Ice.Blobject {
    public boolean ice_invoke(
        byte[] inParams,
        Ice.ByteSeqHolder outParams,
        Ice.Current current)
    {
        // ...
    }
}
```

An instance of `ComputeI` is an Ice object because `Blobject` derives from `Object`, therefore an instance can be added to an object adapter like any other servant.

For the purposes of this discussion, the implementation of `ice_invoke` handles only the `add` operation and raises `OperationNotExistException` for all other operations. In a real implementation, the servant must also be prepared to receive invocations of the following `Object` operations:

Copyright 2016, ZeroC, Inc.

- `string ice_id()`
  Returns the Slice type ID of the servant's most-derived type.
- `StringSeq ice_ids()`
  Returns a sequence of strings representing all of the Slice interfaces supported by the servant, including `"::Ice::Object"`.
- `bool ice_isA(string id)`
  Returns `true` if the servant supports the interface denoted by the given Slice type ID, or `false` otherwise. This operation is invoked by the proxy function `checkedCast`.
- `void ice_ping()`
  Verifies that the object denoted by the identity and facet contained in `Ice::Current` is reachable.

With that in mind, here is our simplified version of `ice_invoke`:

Copyright 2016, ZeroC, Inc.

Java

```
    public boolean ice_invoke(
        byte[] inParams,
        Ice.ByteSeqHolder outParams,
        Ice.Current current)
    {
        if (current.operation.equals("add")) {
            Ice.Communicator communicator = current.adapter.getCommunicator();
            Ice.InputStream in = Ice.Util.createInputStream(communicator,
inParams);
            in.startEncapsulation();
            int x = in.readInt();
            int y = in.readInt();
            in.endEncapsulation();

            Ice.OutputStream out = Ice.Util.createOutputStream(communicator);
            try {
                if (checkOverflow(x, y)) {
                    Calc.Overflow ex = new Calc.Overflow();
                    ex.x = x;
                    ex.y = y;
                    out.startEncapsulation();
                    out.writeException(ex);
                    out.endEncapsulation();
                    outParams.value = out.finished();
                    return false;
                } else {
                    out.startEncapsulation();
                    out.writeInt(x + y);
                    out.endEncapsulation();
                    outParams.value = out.finished();
                    return true;
                }
            } finally {
                out.destroy();
            }
        } else {
            Ice.OperationNotExistException ex =
new Ice.OperationNotExistException();
            ex.id = current.id;
            ex.facet = current.facet;
            ex.operation = current.operation;
            throw ex;
        }
    }
```

If an overflow is detected, the code "raises" the `Calc::Overflow` user exception by calling `writeException` on the output stream and returning `false`, otherwise the return value is encoded and the function returns `true`.

See Also

- Object Adapters

Copyright 2016, ZeroC, Inc.

- Request Contexts
- Streaming Interfaces
- Type IDs
- Object Identity
- Versioning

Copyright 2016, ZeroC, Inc.

Dynamic Invocation and Dispatch in C-Sharp

This page describes the C# mapping for the `ice_invoke` proxy function and the `Blobject` class.

On this page:

- ice_invoke in C#
- Using Streams with ice_invoke in C#
- Subclassing Blobject in C#

`ice_invoke` in C#

The mapping for `ice_invoke` is shown below:

C#

```
namespace Ice
{
    public interface ObjectPrx
    {
        bool ice_invoke(string operation,
                        OperationMode mode,
                        byte[] inParams,
                        out byte[] outParams);
        // ...
    }
}
```

Another overloading of `ice_invoke` (not shown) adds a trailing argument of type `Ice.Context`.

As an example, the code below demonstrates how to invoke the operation `op`, which takes no `in` parameters:

C#

```
Ice.ObjectPrx proxy = ...
try {
    byte[] outParams;
    if (proxy.ice_invoke("op", Ice.OperationMode.Normal, null, outParams)) {
        // Handle success
    } else {
        // Handle user exception
    }
} catch (Ice.LocalException ex) {
    // Handle exception
}
```

As a convenience, the Ice run time accepts a null or empty byte sequence when there are no input parameters and internally translates it into an empty encapsulation. In all other cases, the value for `inParams` must be an encapsulation of the encoded parameters.

Using Streams with `ice_invoke` in C#

The streaming interfaces provide the tools an application needs to dynamically invoke operations with arguments of any Slice type. Consider the following Slice definition:

Copyright 2016, ZeroC, Inc.

Slice

```
module Calc {
    exception Overflow {
        int x;
        int y;
    };
    interface Compute {
        idempotent int add(int x, int y)
            throws Overflow;
    };
};
```

Now let's write a client that dynamically invokes the `add` operation:

C#

```
Ice.ObjectPrx proxy = ...
try {
    Ice.OutputStream outStream = Ice.Util.createOutputStream(communicator);
    outStream.startEncapsulation();
    outStream.writeInt(100); // x
    outStream.writeInt(-1);  // y
    outStream.endEncapsulation();
    byte[] inParams = outStream.finished();

    byte[] outParams;
    if (proxy.ice_invoke("add", Ice.OperationMode.Idempotent,
inParams, out outParams)) {
        // Handle success
        Ice.InputStream inStream =
Ice.Util.createInputStream(communicator, outParams);
        inStream.startEncapsulation();
        int result = inStream.readInt();
        inStream.endEncapsulation();
        System.Diagnostics.Debug.Assert(result == 99);
    } else {
        // Handle user exception
    }
} catch (Ice.LocalException ex) {
    // Handle exception
}
```

You can see here that the input and output parameters are enclosed in encapsulations.

We neglected to handle the case of a user exception in this example, so let's implement that now. We assume that we have compiled our program with the Slice-generated code, therefore we can call `throwException` on the input stream and catch `Overflow` directly:

Copyright 2016, ZeroC, Inc.

C#

```
    if (proxy.ice_invoke("add", Ice.OperationMode.Idempotent,
inParams, out outParams)) {
        // Handle success
        ...
    } else {
        // Handle user exception
        Ice.InputStream inStream =
Ice.Util.createInputStream(communicator, outParams);
        try {
            inStream.startEncapsulation();
            inStream.throwException();
        } catch (Calc.Overflow ex) {
            System.Console.WriteLine("overflow while adding " +
                                     ex.x + " and " + ex.y);
        } catch (Ice.UserException) {
            // Handle unexpected user exception
        }
    }
```

This is obviously a contrived example: if the Slice-generated code is available, why bother using dynamic dispatch? In the absence of Slice-generated code, the caller would need to manually unmarshal the user exception, which is outside the scope of this book.

As a defensive measure, the code traps `Ice.UserException`. This could be raised if the Slice definition of `add` is modified to include another user exception but this segment of code did not get updated accordingly.

Subclassing `Blobject` in C#

Implementing the dynamic dispatch model requires writing a subclass of `Ice.Blobject`. We continue using the `Compute` interface to demonstrate a `Blobject` implementation:

C#

```
public class ComputeI : Ice.Blobject {
    public bool ice_invoke(byte[] inParams, out byte[] outParams,
Ice.Current current);
    {
        ...
    }
}
```

An instance of `ComputeI` is an Ice object because `Blobject` derives from `Object`, therefore an instance can be added to an object adapter like any other servant.

For the purposes of this discussion, the implementation of `ice_invoke` handles only the `add` operation and raises `OperationNotExistException` for all other operations. In a real implementation, the servant must also be prepared to receive invocations of the following `Object` operations:

- `string ice_id()`
  Returns the Slice type ID of the servant's most-derived type.

Copyright 2016, ZeroC, Inc.

- `StringSeq ice_ids()`
  Returns a sequence of strings representing all of the Slice interfaces supported by the servant, including `"::Ice::Object"`.
- `bool ice_isA(string id)`
  Returns `true` if the servant supports the interface denoted by the given Slice type ID, or `false` otherwise. This operation is invoked by the proxy function `checkedCast`.
- `void ice_ping()`
  Verifies that the object denoted by the identity and facet contained in `Ice::Current` is reachable.

With that in mind, here is our simplified version of `ice_invoke`:

Copyright 2016, ZeroC, Inc.

C#

```
    public bool ice_invoke(byte[] inParams, out byte[] outParams,
Ice.Current current);
    {
        if (current.operation.Equals("add")) {
            Ice.Communicator communicator = current.adapter.getCommunicator();
            Ice.InputStream inStream = Ice.Util.createInputStream(communicator,
inParams);
            inStream.startEncapsulation();
            int x = inStream.readInt();
            int y = inStream.readInt();
            inStream.endEncapsulation();

            Ice.OutputStream outStream = Ice.Util.createOutputStream(communicator);
            try {
                if (checkOverflow(x, y)) {
                    Calc.Overflow ex = new Calc.Overflow();
                    ex.x = x;
                    ex.y = y;
                    outStream.startEncapsulation();
                    outStream.writeException(ex);
                    outStream.endEncapsulation();
                    outParams = outStream.finished();
                    return false;
                } else {
                    outStream.startEncapsulation();
                    outStream.writeInt(x + y);
                    outStream.endEncapsulation();
                    outParams = outStream.finished();
                    return true;
                }
            } finally {
                outStream.destroy();
            }
        } else {
            Ice.OperationNotExistException ex =
new Ice.OperationNotExistException();
            ex.id = current.id;
            ex.facet = current.facet;
            ex.operation = current.operation;
            throw ex;
        }
    }
```

If an overflow is detected, the code "raises" the `Calc::Overflow` user exception by calling `writeException` on the output stream and returning `false`, otherwise the return value is encoded and the function returns `true`.

See Also

- Object Adapters
- Request Contexts
- Streaming Interfaces
- Type IDs

Copyright 2016, ZeroC, Inc.

- Object Identity
- Versioning

Copyright 2016, ZeroC, Inc.

Asynchronous Dynamic Invocation and Dispatch

Ice also provides asynchronous support for the dynamic invocation and dispatch models. The mappings for the `ice_invoke` proxy function and the `Blobject` class adhere to the normal asynchronous mapping rules.

Topics

- Asynchronous Dynamic Invocation and Dispatch in C++
- Asynchronous Dynamic Invocation and Dispatch in Java
- Asynchronous Dynamic Invocation and Dispatch in C-Sharp

See Also

- Dynamic Invocation and Dispatch

Copyright 2016, ZeroC, Inc.

Asynchronous Dynamic Invocation and Dispatch in C++

This page describes the asynchronous C++ mapping for the `ice_invoke` proxy function and the `Blobject` class.

On this page:

- Calling ice_invoke Asynchronously in C++
  - Basic Asynchronous Mapping for ice_invoke in C++
  - Generic Asynchronous Callback Mapping for ice_invoke in C++
  - Type-Safe Asynchronous Callback Mapping for ice_invoke in C++
- Subclassing BlobjectAsync in C++
- Subclassing BlobjectArrayAsync in C++

Calling `ice_invoke` Asynchronously in C++

The asynchronous mapping for `ice_invoke` resembles that of the static AMI mapping. Multiple overloadings are provided to support the use of callbacks, request contexts, and zero-copy semantics. The return value and the parameters `operation`, `mode`, and `inParams` have the same semantics as for the synchronous version of `ice_invoke`.

Basic Asynchronous Mapping for `ice_invoke` in C++

The basic mapping is shown below:

C++

```
Ice::AsyncResultPtr begin_ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::vector<Ice::Byte>& inParams);

Ice::AsyncResultPtr begin_ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::vector<Ice::Byte>& inParams,
    const Ice::Context& __ctx);

Ice::AsyncResultPtr begin_ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::pair<const Ice::Byte*, const Ice::Byte*>& inParams);

Ice::AsyncResultPtr begin_ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::pair<const Ice::Byte*, const Ice::Byte*>& inParams,
    const Ice::Context& __ctx,
    const Ice::LocalObjectPtr& __cookie = 0);

bool end_ice_invoke(std::vector<Ice::Byte>&, const Ice::AsyncResultPtr&);
```

User exceptions are handled differently than for static asynchronous invocations. Calling `end_ice_invoke` can raise system exceptions but never raises user exceptions. Instead, the boolean return value of `end_ice_invoke` indicates whether the operation completed successfully (true) or raised a user exception (false). If the return value is true, the byte sequence contains an encapsulation of the results; otherwise, the byte sequence contains an encapsulation of the user exception.

Generic Asynchronous Callback Mapping for `ice_invoke` in C++

Copyright 2016, ZeroC, Inc.

The generic callback API is also available:

C++

```
Ice::AsyncResultPtr begin_ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::vector<Ice::Byte>& inParams,
    const Ice::CallbackPtr& __del,
    const Ice::LocalObjectPtr& __cookie = 0);

Ice::AsyncResultPtr begin_ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::vector<Ice::Byte>& inParams,
    const Ice::Context& __ctx,
    const Ice::CallbackPtr& __del,
    const Ice::LocalObjectPtr& __cookie = 0);

Ice::AsyncResultPtr begin_ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::pair<const Ice::Byte*, const Ice::Byte*>& inParams,
    const Ice::CallbackPtr& __del,
    const Ice::LocalObjectPtr& __cookie = 0);

Ice::AsyncResultPtr begin_ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::pair<const Ice::Byte*, const Ice::Byte*>& inParams,
    const Ice::Context& __ctx,
    const Ice::CallbackPtr& __del,
    const Ice::LocalObjectPtr& __cookie = 0);
```

Use the `Ice::newCallback` function to create callback objects, as shown in the static AMI mapping.

Type-Safe Asynchronous Callback Mapping for `ice_invoke` in C++

The type-safe callback API looks as follows:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::AsyncResultPtr begin_ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::vector<Ice::Byte>& inParams,
    const Ice::Callback_Object_ice_invokePtr& __del,
    const Ice::LocalObjectPtr& __cookie = 0);

Ice::AsyncResultPtr begin_ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::vector<Ice::Byte>& inParams,
    const Ice::Context& __ctx,
    const Ice::Callback_Object_ice_invokePtr& __del,
    const Ice::LocalObjectPtr& __cookie = 0);

Ice::AsyncResultPtr begin_ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::pair<const Ice::Byte*, const Ice::Byte*>& inParams,
    const Ice::Callback_Object_ice_invokePtr& __del,
    const Ice::LocalObjectPtr& __cookie = 0);

Ice::AsyncResultPtr begin_ice_invoke(
    const std::string& operation,
    Ice::OperationMode mode,
    const std::pair<const Ice::Byte*, const Ice::Byte*>& inParams,
    const Ice::Context& __ctx,
    const Ice::Callback_Object_ice_invokePtr& __del,
    const Ice::LocalObjectPtr& __cookie = 0);
```

Several overloadings of `Ice::newCallback_Object_ice_invoke` allow you to create callback objects. Versions are provided to support zero-copy semantics for the byte sequence containing the operation's results, as well as a cookie value whose type is inferred and represented here by the `CT` symbol:

Copyright 2016, ZeroC, Inc.

C++

```
template<class T> Callback_Object_ice_invokePtr
newCallback_Object_ice_invoke(const IceUtil::Handle<T>& instance,
                              void (T::*cb)(bool, const std::vector<Ice::Byte>&),
                              void (T::*excb)(const ::Ice::Exception&),
                              void (T::*sentcb)(bool) = 0);

template<class T> Callback_Object_ice_invokePtr
newCallback_Object_ice_invoke(const IceUtil::Handle<T>& instance,
                              void (T::*cb)(bool, const std::pair<const Byte*,
const Byte*>&),
                              void (T::*excb)(const ::Ice::Exception&),
                              void (T::*sentcb)(bool) = 0);

template<class T, typename CT> Callback_Object_ice_invokePtr
newCallback_Object_ice_invoke(const IceUtil::Handle<T>& instance,
                              void (T::*cb)(bool, const std::vector<Ice::Byte>&,
const CT&),
                              void (T::*excb)(const ::Ice::Exception&, const CT&),
                              void (T::*sentcb)(bool, const CT&) = 0);

template<class T, typename CT> Callback_Object_ice_invokePtr
newCallback_Object_ice_invoke(const IceUtil::Handle<T>& instance,
                              void (T::*cb)(bool, const std::pair<const Byte*,
const Byte*>&,
                                            const CT&),
                              void (T::*excb)(const ::Ice::Exception&, const CT&),
                              void (T::*sentcb)(bool, const CT&) = 0);

template<class T> Callback_Object_ice_invokePtr
newCallback_Object_ice_invoke(const IceUtil::Handle<T>& instance,
                              void (T::*excb)(const ::Ice::Exception&),
                              void (T::*sentcb)(bool) = 0);

template<class T, typename CT> Callback_Object_ice_invokePtr
newCallback_Object_ice_invoke(const IceUtil::Handle<T>& instance,
                              void (T::*excb)(const ::Ice::Exception&, const CT&),
                              void (T::*sentcb)(bool, const CT&) = 0);
```

Subclassing `BlobjectAsync` in C++

`BlobjectAsync` is the name of the asynchronous counterpart to `Blobject`:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Ice {
    class BlobjectAsync : virtual public Ice::Object {
    public:
        virtual void ice_invoke_async(
            const AMD_Object_ice_invokePtr& cb,
            const std::vector<Ice::Byte>& inParams,
            const Ice::Current& current) = 0;
    };
}
```

To implement asynchronous dynamic dispatch, a server must subclass `BlobjectAsync` and override `ice_invoke_async`.

The first argument to the servant's member function is a callback object of type `Ice::AMD_Object_ice_invoke`, shown here:

C++

```
namespace Ice {
    class AMD_Object_ice_invoke : ... {
    public:
        virtual void ice_response(
            bool result,
            const std::vector<Ice::Byte>& outParams) = 0;

        virtual void ice_response(
            bool result,
            const std::pair<const Ice::Byte*, const Ice::Byte*>& outParams) = 0;

        virtual void ice_exception(const std::exception&) = 0;
        virtual void ice_exception() = 0;
    };
}
```

Upon a successful invocation, the servant must invoke one of the `ice_response` methods on the callback object, passing `true` as the first argument and encoding the encapsulated operation results into `outParams`. To report a user exception, the servant invokes `ice_response` with `false` as the first argument and the encapsulated form of the exception in `outParams`. Note that the second overloading of `ice_response` uses the array mapping (see the next section below).

`ice_exception` has several overloadings. Note however that in the dynamic dispatch model, the `ice_exception` function must not be used to report user exceptions; doing so results in the caller receiving `UnknownUserException`.

Subclassing `BlobjectArrayAsync` in C++

To implement an asynchronous Blobject servant that uses the array mapping, derive your implementation class from `Ice::BlobjectArrayAsync` and override the `ice_invoke_async` function:

Copyright 2016, ZeroC, Inc.

C++

```
class BlobjectArrayAsync : virtual public Ice::Object {
public:
    virtual void ice_invoke_async(
        const AMD_Object_ice_invokePtr& cb,
        const std::pair<const Ice::Byte*, const Ice::Byte*>& in,
        const Ice::Current& current) = 0;
};
```

As shown in the previous section, the AMD callback class provides an overloaded `ice_response` method that supports the array mapping for the encoded `out` parameter blob.

The discussion of sequences provides more information on the array mapping.

See Also

- Asynchronous Method Invocation (AMI) in C++
- Request Contexts
- C++ Mapping for Sequences
- Dynamic Invocation and Dispatch in C++

Copyright 2016, ZeroC, Inc.

Asynchronous Dynamic Invocation and Dispatch in Java

This page describes the asynchronous Java mapping for the `ice_invoke` proxy function and the `Blobject` class.

On this page:

- Calling ice_invoke Asynchronously in Java
  - Basic Asynchronous Mapping for ice_invoke in Java
  - Generic Asynchronous Callback Mapping for ice_invoke in Java
  - Type-Safe Asynchronous Callback Mapping for ice_invoke in Java
- Subclassing BlobjectAsync in Java

Calling `ice_invoke` Asynchronously in Java

The asynchronous mapping for `ice_invoke` resembles that of the static AMI mapping. Multiple overloadings are provided to support the use of callbacks and request contexts. The return value and the parameters `operation`, `mode`, and `inParams` have the same semantics as for the synchronous version of `ice_invoke`.

Basic Asynchronous Mapping for `ice_invoke` in Java

The basic mapping is shown below:

Java

```
Ice.AsyncResult begin_ice_invoke(
    String operation,
    Ice.OperationMode mode,
    byte[] inParams);

Ice.AsyncResult begin_ice_invoke(
    String operation,
    Ice.OperationMode mode,
    byte[] inParams,
    java.util.Map<String, String> __context);

boolean end_ice_invoke(Ice.ByteSeqHolder outParams, Ice.AsyncResult __result);
```

User exceptions are handled differently than for static asynchronous invocations. Calling `end_ice_invoke` can raise system exceptions but never raises user exceptions. Instead, the boolean return value of `end_ice_invoke` indicates whether the operation completed successfully (true) or raised a user exception (false). If the return value is true, the byte sequence contains an encapsulation of the results; otherwise, the byte sequence contains an encapsulation of the user exception.

Generic Asynchronous Callback Mapping for `ice_invoke` in Java

The generic callback API is also available:

Copyright 2016, ZeroC, Inc.

Java

```
Ice.AsyncResult begin_ice_invoke(
    String operation,
    Ice.OperationMode mode,
    byte[] inParams,
    Ice.Callback __cb);

Ice.AsyncResult begin_ice_invoke(
    String operation,
    Ice.OperationMode mode,
    byte[] inParams,
    java.util.Map<String, String> __context,
    Ice.Callback __cb);
```

Refer to the static AMI mapping for an example of subclassing `Ice.Callback`.

Type-Safe Asynchronous Callback Mapping for `ice_invoke` in Java

The type-safe callback API looks as follows:

Java

```
Ice.AsyncResult begin_ice_invoke(
    String operation,
    Ice.OperationMode mode,
    byte[] inParams,
    Ice.Callback_Object_ice_invoke __cb);

Ice.AsyncResult begin_ice_invoke(
    String operation,
    Ice.OperationMode mode,
    byte[] inParams,
    java.util.Map<String, String> __context,
    Ice.Callback_Object_ice_invoke __cb);
```

Callers must supply a subclass of `Ice.Callback_Object_ice_invoke`:

Copyright 2016, ZeroC, Inc.

Java

```
package Ice;

public abstract class Callback_Object_ice_invoke extends ...
{
    public abstract void response(boolean __ret, byte[] outParams);

    public abstract void exception(LocalException __ex);
}
```

The boolean argument to `response` indicates whether the operation completed successfully (true) or raised a user exception (false). If the return value is true, the byte sequence contains an encapsulation of the results; otherwise, the byte sequence contains an encapsulation of the user exception.

Subclassing `BlobjectAsync` in Java

`BlobjectAsync` is the name of the asynchronous counterpart to `Blobject`:

Java

```
package Ice;

public abstract class BlobjectAsync extends Ice.ObjectImpl {
    public abstract void ice_invoke_async(
        Ice.AMD_Object_ice_invoke cb,
        byte[] inParams,
        Ice.Current current);

    // ...
}
```

To implement asynchronous dynamic dispatch, a server must subclass `BlobjectAsync` and override `ice_invoke_async`.

As with any other asynchronous operation, the first argument to the servant's member function is always a callback object. In this case, the callback object is of type `Ice.AMD_Object_ice_invoke`, shown here:

Java

```
package Ice;

public interface AMD_Object_ice_invoke {
    void ice_response(boolean result, byte[] outParams);
    void ice_exception(java.lang.Exception ex);
}
```

Upon a successful invocation, the servant must invoke `ice_response` on the callback object, passing `true` as the first argument and encoding the encapsulated operation results into `outParams`. To report a user exception, the servant invokes `ice_response` with `false` a

Copyright 2016, ZeroC, Inc.

s the first argument and the encapsulated form of the exception in `outParams`.

In the dynamic dispatch model, the `ice_exception` function must not be used to report user exceptions; doing so results in the caller receiving `UnknownUserException`.

See Also

- Asynchronous Method Invocation (AMI) in Java
- Request Contexts
- Dynamic Invocation and Dispatch Overview

Copyright 2016, ZeroC, Inc.

Asynchronous Dynamic Invocation and Dispatch in C-Sharp

This page describes the asynchronous C# mapping for the `ice_invoke` proxy function and the `Blobject` class.

On this page:

- Calling ice_invoke Asynchronously in C#
  - Basic Asynchronous Mapping for ice_invoke in C#
  - Generic Asynchronous Callback Mapping for ice_invoke in C#
  - Type-Safe Asynchronous Callback Mapping for ice_invoke in C#
- Subclassing BlobjectAsync in C#

Calling `ice_invoke` Asynchronously in C#

The asynchronous mapping for `ice_invoke` resembles that of the static AMI mapping. Multiple overloadings are provided to support the use of callbacks and request contexts. The return value and the parameters `operation`, `mode`, and `inParams` have the same semantics as for the synchronous version of `ice_invoke`.

Basic Asynchronous Mapping for `ice_invoke` in C#

The basic mapping is shown below:

C#

```
Ice.AsyncResult<Ice.Callback_Object_ice_invoke>
begin_ice_invoke(
    string operation,
    Ice.OperationMode mode,
    byte[] inParams);

Ice.AsyncResult<Callback_Object_ice_invoke>
begin_ice_invoke(
    string operation,
    Ice.OperationMode mode,
    byte[] inParams,
    Dictionary<string, string> context__);

bool end_ice_invoke(out byte[] outParams, AsyncResult r__);
```

User exceptions are handled differently than for static asynchronous invocations. Calling `end_ice_invoke` can raise system exceptions but never raises user exceptions. Instead, the boolean return value of `end_ice_invoke` indicates whether the operation completed successfully (true) or raised a user exception (false). If the return value is true, the byte sequence contains an encapsulation of the results; otherwise, the byte sequence contains an encapsulation of the user exception.

Generic Asynchronous Callback Mapping for `ice_invoke` in C#

The generic callback API is also available:

Copyright 2016, ZeroC, Inc.

C#

```
Ice.AsyncResult begin_ice_invoke(
    string operation,
    Ice.OperationMode mode,
    byte[] inParams,
    Ice.AsyncCallback cb__,
    object cookie__);

Ice.AsyncResult begin_ice_invoke(
    string operation,
    Ice.OperationMode mode,
    byte[] inParams,
    Dictionary<string, string> context__,
    Ice.AsyncCallback cb__,
    object cookie__);
```

Refer to the static AMI mapping for a callback example.

Type-Safe Asynchronous Callback Mapping for `ice_invoke` in C#

For the type-safe callback API, you register callbacks on the `AsyncResult` object just as in the static AMI mapping:

C#

```
public class MyCallback
{
    public void responseCB(bool ret, byte[] results)
    {
        if(ret)
            System.Console.Out.WriteLine("Success");
        else
            System.Console.Out.WriteLine("User exception");
    }

    public void failureCB(Ice.Exception ex)
    {
        System.Console.Err.WriteLine("Exception is: " + ex);
    }
}

...

Ice.AsyncResult<Ice.Callback_Object_ice_invoke> r = proxy.begin_ice_invoke(...);
MyCallback cb = new MyCallback();
r.whenCompleted(cb.responseCB, cb.failureCB);
```

The caller invokes `whenCompleted` on the `AsyncResult` object and supplies delegates to handle response and failure. The response delegate must match the signature of `Ice.Callback_Object_ice_invoke`:

Copyright 2016, ZeroC, Inc.

C#

```
public delegate void Callback_Object_ice_invoke(bool ret__, byte[] outParams);
```

Subclassing `BlobjectAsync` in C#

`BlobjectAsync` is the name of the asynchronous counterpart to `Blobject`:

C#

```
public abstract class BlobjectAsync : Ice.ObjectImpl
{
    public abstract void ice_invoke_async(
        AMD_Object_ice_invoke cb,
        byte[] inParams,
        Current current);
}
```

To implement asynchronous dynamic dispatch, a server must subclass `BlobjectAsync` and override `ice_invoke_async`.

The first argument to the servant's member function is a callback object of type `Ice.AMD_Object_ice_invoke`, shown here:

C#

```
namespace Ice
{
    public interface AMD_Object_ice_invoke
    {
        void ice_response(bool ok, byte[] outParams);
        void ice_exception(System.Exception ex);
    }
}
```

Upon a successful invocation, the servant must invoke `ice_response` on the callback object, passing `true` as the first argument and encoding the encapsulated operation results into `outParams`. To report a user exception, the servant invokes `ice_response` with `false` a s the first argument and the encapsulated form of the exception in `outParams`.

In the dynamic dispatch model, the ice_exception function must not be used to report user exceptions; doing so results in the caller receiving `UnknownUserException`.

See Also

- Asynchronous Method Invocation (AMI) in C-Sharp
- Request Contexts
- Dynamic Invocation and Dispatch Overview

Copyright 2016, ZeroC, Inc.

Facets

Facets provide a general-purpose mechanism for non-intrusively extending the type system of an application, by loosely coupling new type instances to existing ones. This shifts the type selection process from compile to run time and implements a form of late binding. This is particularly useful for versioning an application.

On this page:

- Ice Objects as Collections of Facets
- Server-Side Facet Operations
- Client-Side Facet Operations
- Facet Exception Semantics

Ice Objects as Collections of Facets

Up to this point, we have presented an Ice object as a single conceptual entity, that is, as an object with a single most-derived interface and a single identity, with the object being implemented by a single servant. However, an Ice object is more correctly viewed as a collection of one or more sub-objects known as facets, as shown below:

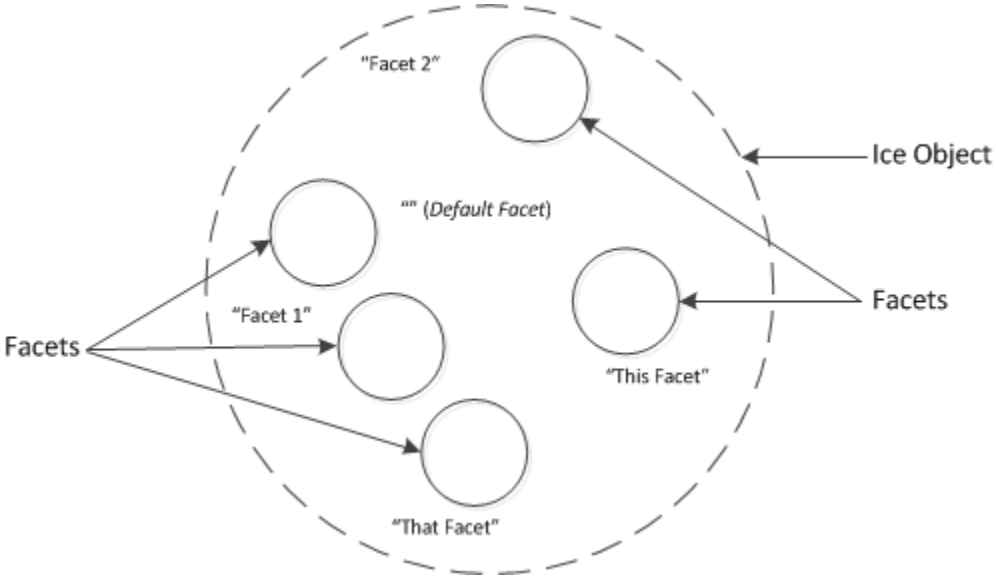

An Ice object with five facets sharing a single object identity.

The diagram above shows a single Ice object with five facets. Each facet has a name, known as the *facet name*. Within a single Ice object, all facets must have unique names. Facet names are arbitrary strings that are assigned by the server that implements an Ice object. A facet with an empty facet name is legal and known as the *default facet*. Unless you arrange otherwise, an Ice object has a single default facet; by default, operations that involve Ice objects and servants operate on the default facet.

Note that all the facets of an Ice object share the same single identity, but have different facet names. Recall the definition of `Ice::Current` once more:

Copyright 2016, ZeroC, Inc.

Slice

```
module Ice {
    local dictionary<string, string> Context;

    enum OperationMode { Normal, \Nonmutating, \Idempotent };

    local struct Current {
        ObjectAdapter   adapter;
        Identity        id;
        string          facet;
        string          operation;
        OperationMode   mode;
        Context         ctx;
        int             requestId;
        EncodingVersion encoding;
    };
};
```

By definition, if two facets have the same `id` field, they are part of the same Ice object. Also by definition, if two facets have the same `id` fiel d, their `facet` fields have different values.

Even though Ice objects usually consist of just the default facet, it is entirely legal for an Ice object to consist of facets that all have non-empty names (that is, it is legal for an Ice object not to have a default facet).

Each facet has a single most-derived interface. There is no need for the interface types of the facets of an Ice object to be unique. It is legal for two facets of an Ice object to implement the same most-derived interface.

Each facet is implemented by a servant. All the usual implementation techniques for servants are available to implement facets — for example, you can implement a facet using a servant locator. Typically, each facet of an Ice object has a separate servant, although, if two facets of an Ice object have the same type, they can also be implemented by a single servant (for example, using a default servant).

Server-Side Facet Operations

On the server side, the object adapter offers a number of operations to support facets:

Slice

```
namespace Ice {
  dictionary<string, Object> FacetMap;

  local interface ObjectAdapter {
     Object*  addFacet(Object servant, Identity id, string facet);
     Object*  addFacetWithUUID(Object servant, string facet);
     Object   removeFacet(Identity id, string facet);
     Object   findFacet(Identity id, string facet);

     FacetMap findAllFacets(Identity id);
     FacetMap removeAllFacets(Identity id);
     // ...
  };
};
```


Copyright 2016, ZeroC, Inc.

These operations have the same semantics as the corresponding "normal" operations for servant activation and deactivation (`add`, `addWithUUID`, `remove`, and `find`), but also accept a facet name. The corresponding "normal" operations are simply convenience operations that supply an empty facet name. For example, `remove(id)` is equivalent to `removeFacet(id, "")`, that is, `remove(id)` operates on the default facet.

`findAllFacets` returns a dictionary of *<facet-name, servant>* pairs that contains all the facets for the given identity.

`removeAllFacets` removes all facets for a given identity from the active servant map, that is, it removes the corresponding Ice object entirely. The operation returns a dictionary of *<facet-name, servant>* pairs that contains all the removed facets.

These operations are sufficient for the server to create Ice objects with any number of facets. For example, assume that we have the following Slice definitions:

Slice

```
module Filesystem {
    // ...

    interface File extends Node {
        idempotent Lines read();
        idempotent void write(Lines text) throws GenericError;
    };
};

module FilesystemExtensions {
    // ...

    class DateTime extends TimeOfDay {
        // ...
    };

    struct Times {
        DateTime createdDate;
        DateTime accessedDate;
        DateTime modifiedDate;
    };

    interface Stat {
        idempotent Times getTimes();
    };
};
```

Here, we have a `File` interface that provides operations to read and write a file, and a `Stat` interface that provides access to the file creation, access, and modification time. (Note that the `Stat` interface is defined in a different module and could also be defined in a different source file.) If the server wants to create an Ice object that contains a `File` instance as the default facet and a `Stat` instance that provides access to the time details of the file, it could do so as follows:

Copyright 2016, ZeroC, Inc.

C++

```
// Create a File instance.
//
Filesystem::FilePtr file = new FileI;

// Create a Stat instance.
//
FilesystemExctensions::DateTimePtr dt = new FilesystemExtensions::DateTime;
FilesystemExtensions::Times times;
times.createdDate = dt;
times.accessedDate = dt;
times.modifiedDate = dt;
FilesystemExtensions::StatPtr stat = new StatI(times);

// Register the File instance as the default facet.
//
Filesystem::FilePrx filePrx = myAdapter->addWithUUID(file);

// Register the Stat instance as a facet with name "Stat".
//
myAdapter->addFacet(stat, filePrx->ice_getIdentity(), "Stat");
```

The first few lines simply create and initialize a `FileI` and `StatI` instance. (The details of this do not matter here.) All the action is in the last two statements:

C++

```
Filesystem::FilePrx filePrx = myAdapter->addWithUUID(file);
myAdapter->addFacet(stat, filePrx->ice_getIdentity(), "Stat");
```

This registers the `FileI` instance with the object adapter as usual. (In this case, we let the Ice run time generate a UUID as the object identity.) Because we are calling `addWithUUID` (as opposed to `addFacetWithUUID`), the instance becomes the default facet.

The second line adds a facet to the instance with the facet name `Stat`. Note that we call `ice_getIdentity` on the `File` proxy to pass an object identity to `addFacet`. This guarantees that the two facets share the same object identity.

Note that, in general, it is a good idea to use `ice_getIdentity` to obtain the identity of an existing facet when adding a new facet. That way, it is guaranteed that the facets share the same identity. (If you accidentally pass a different identity to `addFacet`, you will not add a facet to an existing Ice object, but instead register a new Ice object; using `ice_getIdentity` makes this mistake impossible.)

Client-Side Facet Operations

On the client side, which facet a request is addressed to is implicit in the proxy that is used to send the request. For an application that does not use facets, the facet name is always empty so, by default, requests are sent to the default facet.

The client can use a `checkedCast` to obtain a proxy for a particular facet. For example, assume that the client obtains a proxy to a `File` ins tance as shown above. The client can cast between the `File` facet and the `Stat` facet (and back) as follows:

Copyright 2016, ZeroC, Inc.

C++

```
// Get a File proxy.
//
Filesystem::FilePrx file = ...;

// Get the Stat facet.
//
FilesystemExtensions::StatPrx stat =
    FilesystemExtensions::StatPrx::checkedCast(file, "Stat");

// Go back from the Stat facet to the File facet.
//
Filesystem::FilePrx file2 = Filesystem::FilePrx::checkedCast(stat, "");

assert(file2 == file); // The two proxies are identical.
```

This example illustrates that, given any facet of an Ice object, the client can navigate to any other facet by using a `checkedCast` with the facet name.

If an Ice object does not provide the specified facet, `checkedCast` returns null:

C++

```
FilesystemExtensions::StatPrx stat =
    FilesystemExtensions::StatPrx::checkedCast(file, "Stat");

if (!stat) {
    // No Stat facet on this object, handle error...
} else {
    FilesystemExtensions::Times times = stat->getTimes();

    // Use times struct...
}
```

Note that `checkedCast` also returns a null proxy if a facet exists, but the cast is to the wrong type. For example:

C++

```
// Get a File proxy.
//
Filesystem::FilePrx file = ...;

// Cast to the wrong type.
//
SomeTypePrx prx = SomeTypePrx::checkedCast(file, "Stat");

assert(!prx); // checkedCast returns a null proxy.
```

If you want to distinguish between non-existence of a facet and the facet being of the incorrect type, you can first obtain the facet as type `Ob`

Copyright 2016, ZeroC, Inc.

`ject` and then down-cast to the correct type:

C++

```
// Get a File proxy.
//
Filesystem::FilePrx file = ...;

// Get the facet as type Object.
//
Ice::ObjectPrx obj = Ice::ObjectPrx::checkedCast(file, "Stat");
if (!obj) {
    // No facet with name "Stat" on this Ice object.
} else {
    FilesystemExtensions::StatPrx stat =
        FilesystemExtensions::StatPrx::checkedCast(file);
    if (!stat) {
        // There is a facet with name "Stat", but it is not
        // of type FilesystemExtensions::Stat.
    } else {
        // Use stat...
    }
}
```

This last example also illustrates that

C++

```
StatPrx::checkedCast(prx, "")
```

is *not* the same as

C++

```
StatPrx::checkedCast(prx)
```

The first version explicitly requests a cast to the default facet. This means that the Ice run time first looks for a facet with the empty name and then attempts to down-cast that facet (if it exists) to the type `Stat`.

The second version requests a down-cast that *preserves* whatever facet is currently effective in the proxy. For example, if the `prx` proxy currently holds the facet name "Joe", then (if `prx` points at an object of type `Stat`) the run time returns a proxy of type `StatPrx` that also stores the facet name "Joe".

It follows that, to navigate between facets, you must always use the two-argument version of `checkedCast`, whereas, to down-cast to another type while preserving the facet name, you must always use the single-argument version of `checkedCast`.

You can always check what the current facet of a proxy is by calling `ice_getFacet`:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx obj = ...;

cout << obj->ice_getFacet() << endl; // Print facet name
```

This prints the facet name. (For the default facet, `ice_getFacet` returns the empty string.)

Facet Exception Semantics

The common exceptions `ObjectNotExistException` and `FacetNotExistException` have the following semantics:

- `ObjectNotExistException`
  This exception is raised only if no facets exist at all for a given object identity.
- `FacetNotExistException`
  This exception is raised only if at least one facet exists for a given object identity, but not the specific facet that is the target of an operation invocation.

If you are using servant locators or default servants, you must take care to preserve these semantics. In particular, if you return null from a servant locator's `locate` operation, this appears to the client as an `ObjectNotExistException`. If the object identity for a request is known (that is, there is at least one facet with that identity), but no facet with the specified name exists, you must explicitly throw a `FacetNotExistException` from `locate` instead of simply returning null.

See Also

- Object Identity
- Run-Time Exceptions
- Object Adapters
- Servant Locators
- Default Servants
- The Current Object

Copyright 2016, ZeroC, Inc.

Versioning

Once you have developed and deployed a distributed application, and once the application has been in use for some time, it is likely that you will want to make some changes to the application. For example, you may want to add new functionality to a later version of the application, or you may want to change some existing aspect of the application. Of course, ideally, such changes are accomplished without breaking already deployed software, that is, the changes should be backward compatible. Evolving an application in this way is generally known as *ve rsioning*.

Topics

- Versioning through Incremental Updates
- Versioning with Facets

Copyright 2016, ZeroC, Inc.

Versioning through Incremental Updates

Ice allows you to update your applications and Slice definitions in a backwards-compatible manner - in such a way that the existing applications (using the existing Slice definitions) and the updated applications (using the updated Slice definitions) can easily and safely communicate with each other.

You can safely update your Slice using the techniques described on this page:

- Adding Operations
- Optional Operation Parameters and Class Data Members

Adding Operations

Suppose that we have deployed our file system application and want to add extra functionality to a new version. Specifically, let us assume that the original version (version 1) only provides the basic functionality to use files, but does not provide extra information, such as the modification date or the file size. The question is then, how can we upgrade the existing application with this new functionality? Here is a small excerpt of the original (version 1) Slice definitions once more:

Slice

```
// Version 1

module Filesystem {
    // ...

    interface File extends Node {
        idempotent Lines read();
        idempotent void write(Lines text) throws GenericError;
    };
};
```

Your first attempt at upgrading the application might look as follows:

Copyright 2016, ZeroC, Inc.

Slice

```
// Version 2

module Filesystem {
    // ...

    class DateTime extends TimeOfDay {  // New in version 2
        // ...
    };

    class Times {                       // New in version 2
        DateTime createdDate;
        DateTime accessedDate;
    };

    interface File extends Node {
        idempotent Lines read();
        idempotent void write(Lines text) throws GenericError;

        idempotent Times getTimes();    // New in version 2
    };
};
```

Note that the version 2 definition does not change anything that was present in version 1; instead, it only adds two new types and adds an operation to the `File` interface. Version 1 clients can continue to work with both version 1 and version 2 `File` objects because version 1 clients do not know about the `getTimes` operation and therefore will not call it; version 2 clients, on the other hand, can take advantage of the new functionality. The reason this works is that the Ice protocol invokes an operation by sending the operation name as a string on the wire (rather than using an ordinal number or hash value to identify the operation). Ice guarantees that any future version of the protocol will retain this behavior, so it is safe to add a new operation to an existing interface without recompiling all clients.

However, this approach contains a pitfall: the tacit assumption built into this approach is that no version 2 client will ever use a version 1 object. If the assumption is violated (that is, a version 2 client uses a version 1 object), the version 2 client will receive an `OperationNotExistException` when it invokes the new `getTimes` operation because that operation is supported only by version 2 objects.

Whether you can make this assumption depends on your application. In some cases, it may be possible to ensure that version 2 clients will never access a version 1 object, for example, by simultaneously upgrading all servers from version 1 to version 2, or by taking advantage of application-specific constraints that ensure that version 2 clients only contact version 2 objects. However, for some applications, doing this is impractical.

Note that you could write version 2 clients to catch and react to an `OperationNotExistException` when they invoke the `getTimes` operation: if the operation succeeds, the client is dealing with a version 2 object, and if the operation raises `OperationNotExistsException`, the client is dealing with a version 1 object.

Optional Operation Parameters and Class Data Members

You can safely add one or more optional parameters to an operation without breaking clients or servers that don't know anything of these new parameters.

For example, we could add an optional parameter to the `read` operation n our Filesystem `File` interface optionally limit the number of lines returned:

Copyright 2016, ZeroC, Inc.

Slice

```
interface File extends Node {
    idempotent Lines read(optional (1) int max);
    ...
};
```

A client using the updated Slice definition can set this maximum value, and if its request is received by a new server, the server will receive this max value and should behave accordingly. An old server wouldn't receive this optional parameter, and would continue to behave as before.

You can likewise add optional data members to an existing class without breaking existing applications that use this class. Let's go back to the FileSystem version 2 presented above - say we need to add a third data member to Times, while preserving existing applications that use the existing two-data-members Times. The solution is to add an optional data member, for example:

Slice

```
class Times {
    DateTime createdDate;
    DateTime accessedDate;
    optional (1) DateTime modifiedDate;
};
```

Older applications that know nothing about `modifiedDate` will continue to receive, create, send (etc.) `Times` instances as before (without a `modifiedDate` data member), while newer applications can set and retrieve the optional `modifiedDate` member.

See Also

- Slice for a Simple File System
- Optional Values

Copyright 2016, ZeroC, Inc.

Versioning with Facets

In the most general sense, facets provide a mechanism for implementing multiple interfaces for a single object. The key point is that, to add a new interface to an object, none of the existing definitions have to be touched, so no compatibility issues can arise. More importantly, the decision as to which facet to use is made at run time instead of at compile time.

Used judiciously, facets can handle versioning requirements when the incremental updates approach described earlier is no longer sufficient.

On this page:

- Facet Selection
- Behavioral Versioning
- Facets Design Considerations

Facet Selection

Given that we have decided to extend an application with facets, we have to deal with the question of how clients select the correct facet. The answer typically involves an explicit selection of a facet sometime during client start-up. For example, in our file system application, clients always begin their interactions with the file system by creating a proxy to the root directory. Let us assume that our versioning requirements have led to version 1 and version 2 definitions of directories as follows:

Slice

```
module Filesystem { // Original version
    // ...

    interface Directory extends Node {
        idempotent NodeSeq list();
        // ...
    };
};

module FilesystemV2 {
    // ...

    enum NodeType { Directory, File };

    class NodeDetails {
        NodeType type;
        string name;
        DateTime createdTime;
        DateTime accessedTime;
        DateTime modifiedTime;
        // ...
    };

    interface Directory extends Filesystem::Node {
       idempotent NodeDetailsSeq list();
       // ...
    };
};
```

In this case, the semantics of the `list` operation have changed in version 2. A version 1 client uses the following code to obtain a proxy to the root directory:

Copyright 2016, ZeroC, Inc.

C++

```
// Create a proxy for the root directory
//
Ice::ObjectPrx base = communicator()->stringToProxy("RootDir:default -p 10000");
if (!base)
    throw "Could not create proxy";

// Down-cast the proxy to a Directory proxy
//
Filesystem::DirectoryPrx rootDir = Filesystem::DirectoryPrx::checkedCast(base);
if (!rootDir)
    throw "Invalid proxy";
```

For a version 2 client, the bootstrap code is almost identical — instead of down-casting to `Filesystem::Directory`, the client selects the "V2" facet during the down-cast to the type `FilesystemV2::Directory`:

C++

```
// Create a proxy for the root directory
//
Ice::ObjectPrx base = communicator()->stringToProxy("RootDir:default -p 10000");
if (!base)
    throw "Could not create proxy";

// Down-cast the proxy to a V2 Directory proxy
//
FilesystemV2::DirectoryPrx rootDir
= FilesystemV2::DirectoryPrx::checkedCast(base, "V2");
if (!rootDir)
    throw "Invalid proxy";
```

Of course, we can also create a client that can deal with both version 1 and version 2 directories: if the down-cast to version 2 fails, the client is dealing with a version 1 server and can adjust its behavior accordingly.

Behavioral Versioning

On occasion, versioning requires changes in behavior that are not manifest in the interface of the system. For example, we may have an operation that performs some work, such as:

Slice

```
interface Foo {
    void doSomething();
};
```

The same operation on the same interface exists in both versions, but the *behavior* of `doSomething` in version 2 differs from that in version 1. The question is, how do we best deal with such behavioral changes?

Of course, one option is to simply create a version 2 facet and to carry that facet alongside the original version 1 facet. For example:

Copyright 2016, ZeroC, Inc.

Slice

```
module V2 {

    interface Foo {    // V2 facet
        void doSomething();
    };
};
```

This works fine, as far as it goes: a version 2 client asks for the "V2" facet and then calls `doSomething` to get the desired effect. Depending on your circumstances, this approach may be entirely reasonable. However, if there are such behavioral changes on several interfaces, the approach leads to a more complex type system because it duplicates each interface with such a change.

A better alternative can be to create two facets of the same type, but have the implementation of those facets differ. With this approach, both facets are of type `::Foo::doSomething`. However, the implementation of `doSomething` checks which facet was used to invoke the request and adjusts its behavior accordingly:

C++

```
void
FooI::doSomething(const Ice::Current& c)
{
    if (c.facet == "V2") {
        // Provide version 2 behavior...
    } else {
        // Provide version 1 behavior...
    }
}
```

This approach avoids creating separate types for the different behaviors, but has the disadvantage that version 1 and version 2 objects are no longer distinguishable to the type system. This can matter if, for example, an operation accepts a `Foo` proxy as a parameter. Let us assume that we also have an interface `FooProcessor` as follows:

Slice

```
interface FooProcessor {
    void processFoo(Foo* w);
};
```

If `FooProcessor` also exists as a version 1 and version 2 facet, we must deal with the question of what should happen if a version 1 `Foo` proxy is passed to a version 2 `processFoo` operation because, at the type level, there is nothing to prevent this from happening.

You have two options for dealing with this situation:

- Define working semantics for mixed-version invocations. In this case, you must come up with sensible system behavior for all possible combinations of versions.
- If some of the combinations are disallowed (such as passing a version 1 `Foo` proxy to a version 2 `processFoo` operation), you can detect the version mismatch in the server by looking at the `Current::facet` member and throwing an exception to indicate a version mismatch. Simultaneously, write your clients to ensure they only pass a permissible version to `processFoo`. Clients can ensure this by checking the facet name of a proxy before passing it to `processFoo` and, if there is a version mismatch, changing either the `Foo` proxy or the `FooProcessor` proxy to a matching facet:

Copyright 2016, ZeroC, Inc.

C++

```
FooPrx foo = ...;               // Get a Foo...
FooProcessorPrx fooP = ...;     // Get a FooProcessor...

string fooFacet = foo->ice_getFacet();
string fooPFacet = fooP->ice_getFacet();
if (fooFacet != fooPFacet) {
    if (fooPFacet == "V2") {
        error("Cannot pass a V1 Foo to a V2 FooProcessor");
    } else {
        // Upgrade FooProcessor from V1 to V2
        fooP = FooProcessorPrx::checkedCast(fooP, "V2");
        if (!fooP) {
            error("FooProcessor does not have a V2 facet");
        } else {
            fooP->processFoo(foo);
        }
    }
}
```

Facets Design Considerations

Facets allow you to add versioning to a system, but they are merely a mechanism, not a solution. You still have to make a decision as to how to version something. For example, eventually you may want to deprecate a previous version's behavior; at that point, you must make a decision how to handle requests for the deprecated version. For behavioral changes, you have to decide whether to use separate interfaces or use facets with the same interface. And, of course, you must have compatibility rules to determine what should happen if, for example, a version 1 object is passed to an operation that implements version 2 behavior. In other words, facets cannot do your thinking for you and are no panacea for the versioning problem.

The biggest advantage of facets is also the biggest drawback: facets delay the decision about the types that are used and their behavior until run time. While this provides a lot of flexibility, it is significantly less type-safe than having explicit types that can be statically checked at compile time: if you have a problem relating to incorrect facet selection, the problem will be visible only at run time and, moreover, will be visible only if you actually execute the code that contains the problem, and execute it with just the right data.

Another danger of facets is the opportunity for abuse. As an extreme example, here is an interface that provides an arbitrary collection of objects of arbitrary type:

Slice

```
interface Collection {};
```

Even though this interface is empty, it can provide access to an unlimited number of objects of arbitrary type in the form of facets. While this example is extreme, it illustrates the design tension that is created by facets: you must decide, for a given versioning problem, how and at what point of the type hierarchy to split off a facet that deals with the changed functionality. The temptation may be to "simply add another facet" and be done with it. However, if you do that, your objects are in danger of being nothing more than loose conglomerates of facets without rhyme or reason, and with little visibility of their relationships in the type system.

In object modeling terms, the relationship among facets is weaker than an *is-a* relationship (because facets are often not type-compatible among each other). On the other hand, the relationship among facets is stronger than a *has-a* relationship (because all facets of an Ice object share the same object identity).

It is probably best to treat the relationship of a facet to its Ice object with the same respect as an inheritance relationship: if you were omniscient and could have designed your system for all current and future versions simultaneously, many of the operations that end up on separate facets would probably have been in the same interface instead. In other words, adding a facet to an Ice object most often implies that the facet has an *is-partly-a* relationship with its Ice object. In particular, if you think about the life cycle of an Ice object and find that, when an Ice object is deleted, all its facets must be deleted, this is a strong indication of a correct design. On the other hand, if you find that,

Copyright 2016, ZeroC, Inc.

at various times during an Ice object's life cycle, it has a varying number of facets of varying types, that is a good indication that you are using facets incorrectly.

Ultimately, the decision comes down to deciding whether the trade-off of static type safety versus dynamic type safety is worth the convenience and backward compatibility. The answer depends strongly on the design of your system and individual requirements, so we can only give broad advice here. Finally, there will be a point where no amount of facet trickery will get past the point when "yet one more version will be the straw that breaks the camel's back." At that point, it is time to stop supporting older versions and to redesign the system.

See Also

- Facets
- The Current Object
- Object Identity

Copyright 2016, ZeroC, Inc.

Administration and Diagnostics

This section presents the facilities that Ice provides to administer, monitor and troubleshoot Ice-based applications.

Topics

- Administrative Facility
- Logger Facility
- Instrumentation Facility

Copyright 2016, ZeroC, Inc.

Administrative Facility

Ice applications often require remote administration, such as when an IceGrid node needs to gracefully deactivate a running server. The Ice run time provides an extensible, centralized facility for exporting administrative functionality.

Topics

- The admin Object
- Creating the admin Object
- Using the admin Object
- The Process Facet
- The Properties Facet
- The Logger Facet
- The Metrics Facet
- Filtering Administrative Facets
- Custom Administrative Facets
- Security Considerations for Administrative Facets

See Also

- Locator Configuration for a Server
- Object Adapters

Copyright 2016, ZeroC, Inc.

The admin Object

On this page:

- Overview of the admin Object
- Facets of the admin Object

Overview of the admin Object

When the Administrative Facility is enabled, you can configure Ice to host an administrative object in the `Ice.Admin` object adapter, or you can programmatically host this object in your own object adapter.

In this manual, we refer to the administrative object as the *`admin` object*.

You can retrieve a proxy to the admin object associated with your communicator, if any, by calling `getAdmin` on this communicator, for example:

C++

```
Ice::ObjectPrx adminProxy = myCommunicator->getAdmin();
```

Facets of the admin Object

An Ice object is actually a collection of sub-objects known as facets whose types are not necessarily related. Although facets are typically used for extending and versioning types, they also allow a group of interfaces with a common purpose to be consolidated into a single Ice object with an established interface for navigation. These qualities make facets an excellent match for the requirements of the administrative facility.

Each facet of the `admin` object represents a distinct administrative capability. The object does not have a default facet (that is, a facet with an empty name). However, the Ice run time implements several built-in facets that it adds to the `admin` object:

- the `Process` facet
- the `Properties` facet
- the Logger facet
- the `Metrics` facet

An application can control which facets are installed with a configuration property. An application can also install its own facets if necessary. Administrative facets are not required to inherit from a common Slice interface.

See Also

- Creating the admin Object
- Object Identity
- Versioning
- The Process Facet
- The Properties Facet
- Filtering Administrative Facets
- Custom Administrative Facets
- Ice.Admin.*

Copyright 2016, ZeroC, Inc.

Creating the admin Object

The administrative facility is disabled by default. To enable it, you must set the property `Ice.Admin.Enabled` to a numeric value greater than 0, or leave `Ice.Admin.Enabled` unset and specify endpoints for the `Ice.Admin` administrative object adapter using the property `Ice.Admin.Endpoints`. When `Ice.Admin.Enabled` is set to 0 or a negative value, the administrative facility is disabled, and `Ice.Admin.Endpoints` is ignored.

When the administrative facility is enabled, and

- The `Ice.Admin.Endpoints` property specifies endpoints for the Ice.Admin object adapter, and
- The `Ice.Admin.DelayCreation` property is not set, or is set to 0 or a negative value

Ice creates automatically the admin object during communicator initialization, and hosts all its enabled facets in the built-in `Ice.Admin` object adapter.

This admin object and its facets are created at the end of communicator initialization, after the initialization of all plugins. Ice gives this object the identity *`instance-name`*/admin, where *instance-name* is the value of the the `Ice.Admin.InstanceName` property. If `Ice.Admin.InstanceName` is not set or empty, Ice generates a UUID for *instance-name*.

If Ice does not create the admin object during communicator initialization as described above, you need to create the admin object after communicator initialization by calling `getAdmin` or `createAdmin` on the communicator:

Slice

```
local interface Communicator {
    // ...
    Object* getAdmin();
    Object* createAdmin(ObjectAdapter adminAdapter, Identity adminId);
    // ...
};
```

`getAdmin` is typically used to create the admin object when both `Ice.Admin.Enpoints` and `Ice.Admin.DelayCreation` are set. The resulting admin object's enabled facets are hosted in the `Ice.Admin` object adapter. The identity of this admin object is `instance-name`/admin, where *instance-name* corresponds to the value of `Ice.Admin.InstanceName` or a UUID, as described above.

`createAdmin` is used to create the admin object and host all its enabled facets in the object adapter of your choice (the `adminAdapter` parameter), with the identity of your choice (the `adminId` parameter). If `adminAdapter` is null and `Ice.Admin.Endpoints` specifies endpoints for the `Ice.Admin` object adapter, `createAdmin` creates and activates the `Ice.Admin` object adapter and hosts all the admin object's facets in that adapter. Note that only one admin object can be created for a given communicator. `createAdmin` will fail and raise an exception if there is already an admin object.

The administrative facility introduces additional security considerations, therefore the endpoints for the object adapter where the admin object's facets are hosted must be chosen with caution.

See Also

- Ice.Admin.*
- IceGrid and the Administrative Facility
- Security Considerations for Administrative Facets
- Using the admin Object

Copyright 2016, ZeroC, Inc.

Using the admin Object

A program can obtain a proxy for its `admin` object by calling the `getAdmin` operation on a communicator:

Slice

```
module Ice {
local interface Communicator {
    // ...
    Object* getAdmin();
};
};
```

This operation returns a null proxy if the administrative facility is disabled or the application has not created the admin object with `createAdmin`. The proxy returned by `getAdmin` cannot be used for invoking operations because it refers to the default facet and, as we mentioned previously, the `admin` object does not support a default facet. A program must first obtain a new version of the proxy that is configured with the name of a particular administrative facet before invoking operations on it. Although it cannot be used for invocations, the original proxy is still useful because it contains the endpoints of the object adapter where the admin object is hosted, and therefore the program may elect to export that proxy to a remote client.

To administer a program remotely, somehow you must obtain a proxy for the program's `admin` object. There are several ways for the administrative client to accomplish this:

- Construct the proxy itself, assuming that it knows the admin object's identity, facets, and endpoints. The format of the stringified proxy is as follows:
  adminId -f *admin-facet*:*admin-endpoints*
  The identity, represented here by *adminId*, is *instance-name*/admin for admin objects created during communicator initialization or with `getAdmin`, where *instance-name* is the value of the `Ice.Admin.InstanceName` property or a UUID if that property is not defined. (Clearly, the use of a UUID makes the proxy much more difficult for a client to construct on its own.) The name of the administrative facet is supplied as the value of the -f option, and the endpoints of the object adapter that hosts the admin object appear last in the proxy.
- Invoke an application-specific interface for retrieving the `admin` object's proxy.
- Use the `getServerAdmin` operation on the `IceGrid::Admin` interface, if the remote program was activated by IceGrid (see IceGrid Server Activation).

Having obtained the proxy, the administrative client must select a facet before invoking any operations. For example, the code below shows how to obtain the configuration properties of the remote program:

C++

```
// C++
Ice::ObjectPrx adminObj = ...;
Ice::PropertiesAdminPrx propAdmin = Ice::PropertiesAdminPrx::checkedCast(adminObj,
"Properties");
Ice::PropertyDict props = propAdmin->getPropertiesForPrefix("");
```

Here we used an overloaded version of `checkedCast` to supply the facet name of interest (`Properties`). We could have selected the facet using the proxy method `ice_facet` instead, as shown below:

Copyright 2016, ZeroC, Inc.

C++

```
// C++
Ice::ObjectPrx adminObj = ...;
Ice::PropertiesAdminPrx propAdmin = Ice::PropertiesAdminPrx::checkedCast(
    adminObj->ice_facet("Properties"));
Ice::PropertyDict props = propAdmin->getPropertiesForPrefix("");
```

This code is functionally equivalent to the first example.

A remote client must also know (or be able to determine) which facets are available in the target server. Typically this information is statically configured in the client, since the client must also know the interface types of any facets that it uses. If an invocation on a facet raises `FacetNotExistException`, the client may have used an incorrect facet name, or the server may have disabled the facet in question.

See Also

- The admin Object
- Creating the admin Object
- Ice.Admin.*
- Proxy and Endpoint Syntax
- Automatic Retries
- Versioning

Copyright 2016, ZeroC, Inc.

The Process Facet

An activation service, such as an IceGrid node, needs a reliable way to gracefully deactivate a server. One approach is to use a platform-specific mechanism, such as POSIX signals. This works well on POSIX platforms when the server is prepared to intercept signals a nd react appropriately. On Windows platforms, it works less reliably for C++ servers, and not at all for Java servers. For these reasons, the `P rocess` facet provides an alternative that is both portable and reliable.

Be aware of the security considerations associated with enabling the `Process` facet.

On this page:

- The Process Interface
- Obtaining the Local Process Facet
- Application Requirements for the Process Facet
- Replacing the Process Facet
- Integrating the Process Facet with an Activation Service

The `Process` Interface

The Slice interface `Ice::Process` allows an activation service to request a graceful shutdown of the program:

Slice

```
module Ice {
interface Process {
    void shutdown();
    void writeMessage(string message, int fd);
};
};
```

When `shutdown` is invoked, the object implementing this interface is expected to initiate the termination of its process. The activation service may expect the program to terminate within a certain period of time, after which it may terminate the program abruptly.

The `writeMessage` operation allows remote clients to print a message to the program's standard output (`fd` == 1) or standard error (`fd` == 2) channels.

Obtaining the Local Process Facet

We already showed how to obtain a proxy for a remote administrative facet, but suppose you want to interact with the facet in your local address space. The code below shows the necessary steps:

C++

```
Ice::ObjectPtr obj = communicator->findAdminFacet("Process");
if(obj) { // May be nil if the facet is not enabled
    Ice::ProcessPtr facet = Ice::ProcessPtr::dynamicCast(obj);
    ...
}
```

As shown here, the facet is registered with the name `Process` and a regular language cast is used to downcast the base object type to the `Process` interface.

Copyright 2016, ZeroC, Inc.

Application Requirements for the Process Facet

The default implementation of the `Process` facet requires cooperation from an application in order to successfully terminate a process. Specifically, the facet invokes `shutdown` on its communicator and assumes that the application uses this event as a signal to commence its termination procedure. For example, an application typically uses a thread (often the main thread) to call the communicator operation `waitForShutdown`, which blocks the calling thread until the communicator is shut down or destroyed. After `waitForShutdown` returns, the calling thread can initiate a graceful shutdown of its process.

Replacing the Process Facet

You can replace the default `Process` facet if your application requires a different scheme for gracefully shutting itself down. To define your own facet, create a servant that implements the `Ice::Process` interface. As an example, the servant definition shown below duplicates the functionality of the default `Process` facet:

C++

```
class ProcessI : public Ice::Process {
public:
    ProcessI(const Ice::CommunicatorPtr& communicator) :
_communicator(communicator)
    {}

    void shutdown(const Ice::Current&)
    {
        _communicator->shutdown();
    }

    void writeMessage(const string& msg, Ice::Int fd, const Ice::Current&)
    {
        if(fd == 1) cout << msg << endl;
        else if(fd == 2) cerr << msg << endl;
    }

private:
    const Ice::CommunicatorPtr _communicator;
};
```

As you can see, the default implementation of `shutdown` simply shuts down the communicator, which initiates an orderly termination of the Ice run time's server-side components and prevents object adapters from dispatching any new requests. You can add your own application-specific behavior to the `shutdown` method to ensure that your program terminates in a timely manner.

A servant must not invoke `destroy` on its communicator while executing a dispatched operation.

To avoid the risk of a race condition, the recommended strategy for replacing the `Process` facet is to delay creation of the administrative facets until after communicator initialization, so that your application has a chance to replace the facet:

```
# Delay admin object creation for admin object hosted in the Ice.Admin object
adapter
Ice.Admin.DelayCreation=1
```

With `Ice.Admin.DelayCreation` enabled, the application can safely remove the default `Process` facet and install its own:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::CommunicatorPtr communicator = ...
communicator->removeAdminFacet("Process");
Ice::ProcessPtr myProcessFacet = new MyProcessFacet(...);
communicator->addAdminFacet(myProcessFacet, "Process");
```

If you host the admin object in the `Ice.Admin` object adapter, the final step is to create the admin object by calling `getAdmin` on the communicator:

C++

```
communicator->getAdmin();
```

And if you host the admin object in your own object adapter, the final set is to create the admin object with `createAdmin`:

C++

```
communicator->createAdmin(myAdapter, myAdminId);
```

Integrating the Process Facet with an Activation Service

If the `Ice.Admin.ServerId` and `Ice.Default.Locator` properties are defined, the Ice run time performs the following steps after creating the admin object:

- Obtains proxies for the `Process` facet and the default locator
- Invokes `getRegistry` on the proxy to obtain a proxy for the locator registry
- Invokes `setServerProcessProxy` on the locator registry and supplies the value of `Ice.Admin.ServerId` along with a proxy for the `Process` facet

The identifier specified by `Ice.Admin.ServerId` must uniquely identify the process within the locator registry.

In the case of IceGrid, IceGrid defines the `Ice.Admin.ServerId` and `Ice.Default.Locator` properties for each deployed server. IceGrid also supplies a value for `Ice.Admin.Endpoints` if neither this property nor Ice.Admin.Enabled are defined by the server.

See Also

- Communicators
- Versioning
- Portable Signal Handling in C++
- Security Considerations for Administrative Facets
- Object Adapters
- Creating the admin Object
- Ice.Admin.*
- IceGrid

Copyright 2016, ZeroC, Inc.

The Properties Facet

On this page:

- The PropertiesAdmin Interface
- Obtaining the Local Properties Facet
- Property Update Notifications
  - Implementing a Property Update Callback in C++
  - Implementing a Property Update Callback in Java
  - Implementing a Property Update Callback in C#

The `PropertiesAdmin` Interface

An administrator may find it useful to be able to view or modify the configuration properties of a remote Ice application. For example, the Ice Grid administrative tools allow you to query and update the properties of active servers. The `Properties` facet supplies this functionality.

The `Ice::PropertiesAdmin` interface provides access to the communicator's configuration properties:

Slice

```
module Ice {
interface PropertiesAdmin {
    string getProperty(string key);
    PropertyDict getPropertiesForPrefix(string prefix);
    void setProperties(PropertyDict newProperties);
};
};
```

The `getProperty` operation retrieves the value of a single property, and the `getPropertiesForPrefix` operation returns a dictionary of properties whose keys match the given prefix. These operations have the same semantics as those in the `Ice::Properties` interface.

The `setProperties` operation merges the entries in `newProperties` with the communicator's existing properties. If an entry in `newProperties` matches the name of an existing property, that property's value is replaced with the new value. If the new value is an empty string, the property is removed. Any existing properties that are not modified or removed by the entries in `newProperties` are retained with their original values. If the `Ice.Trace.Admin.Properties` property is enabled, Ice logs a message if a call to `setProperties` results in any changes to the property set.

Modifying a program's configuration properties at run time may not have an effect on the program. For example, many of Ice's standard configuration properties are read once during communicator initialization, and never again.

Obtaining the Local Properties Facet

We already showed how to obtain a proxy for a remote administrative facet, but suppose you want to interact with the facet in your local address space. The code below shows the necessary steps:

C++

```
Ice::ObjectPtr obj = communicator->findAdminFacet("Properties");
if(obj) { // May be nil if the facet is not enabled
    Ice::PropertiesAdminPtr facet = Ice::PropertiesAdminPtr::dynamicCast(obj);
    ...
}
```


Copyright 2016, ZeroC, Inc.

As shown here, the facet is registered with the name `Properties` and a regular language cast is used to downcast the base object type to the `PropertiesAdmin` interface.

Property Update Notifications

The Ice run time can notify an application whenever its properties change due to invocations of the `setProperties` operation on the `PropertiesAdmin` interface. This section describes the native API for receiving property updates.

Implementing a Property Update Callback in C++

The `Properties` facet object provided by the Ice run time derives from the class `NativePropertiesAdmin`:

C++

```
namespace Ice {
class NativePropertiesAdmin : virtual public IceUtil::Shared {
public:
    virtual void addUpdateCallback(const PropertiesAdminUpdateCallbackPtr& cb);
    virtual void removeUpdateCallback(const PropertiesAdminUpdateCallbackPtr& cb);
};
typedef IceUtil::Handle<NativePropertiesAdmin> NativePropertiesAdminPtr;
}
```

The application must supply an instance of `PropertiesAdminUpdateCallback`:

C++

```
namespace Ice {
class PropertiesAdminUpdateCallback : virtual public Ice::LocalObject {
public:
    virtual void updated(const PropertyDict& changes) = 0;
};
typedef IceUtil::Handle<PropertiesAdminUpdateCallback>
PropertiesAdminUpdateCallbackPtr;
}
```

The `updated` method receives a `dictionary<string, string>` value representing the properties that were added, changed or removed, with removed properties denoted by an entry whose value is an empty string. It is legal for the `updated` implementation to modify the set of callbacks, and Ice ignores any exceptions it might raise.

The following code demonstrates how to register a callback:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPtr obj = communicator->findAdminFacet("Properties");
if(obj) { // May be nil if the facet is not enabled
    Ice::NativePropertiesAdminPtr facet =
Ice::NativePropertiesAdminPtr::dynamicCast(obj);
    Ice::PropertiesAdminUpdateCallbackPtr myCallback = new ...;
    facet->addUpdateCallback(myCallback);
}
```

Implementing a Property Update Callback in Java

The `Properties` facet object provided by the Ice run time derives from the class `NativePropertiesAdmin`:

Java

```
package Ice;

public interface NativePropertiesAdmin {
    void addUpdateCallback(PropertiesAdminUpdateCallback callback);
    void removeUpdateCallback(PropertiesAdminUpdateCallback callback);
}
```

The application must supply an instance of `PropertiesAdminUpdateCallback`:

Java

```
package Ice;

public interface PropertiesAdminUpdateCallback {
    void updated(java.util.Map<String, String> changes);
}
```

The `updated` method receives a `dictionary<string, string>` value representing the properties that were added, changed or removed, with removed properties denoted by an entry whose value is an empty string. It is legal for the `updated` implementation to modify the set of callbacks, and Ice ignores any exceptions it might raise.

The following code demonstrates how to register a callback:

Copyright 2016, ZeroC, Inc.

Java

```
Ice.Object obj = communicator.findAdminFacet("Properties");
if(obj != null) { // May be nil if the facet is not enabled
    Ice.NativePropertiesAdmin facet = (Ice.NativePropertiesAdmin)obj;
    Ice.PropertiesAdminUpdateCallback myCallback = new ...;
    facet.addUpdateCallback(myCallback);
}
```

Implementing a Property Update Callback in C#

The `Properties` facet object provided by the Ice run time derives from the class `NativePropertiesAdmin`:

C#

```
namespace Ice {
    public interface NativePropertiesAdmin {
        void addUpdateCallback(PropertiesAdminUpdateCallback callback);
        void removeUpdateCallback(PropertiesAdminUpdateCallback callback);
    }
}
```

The application must supply an instance of `PropertiesAdminUpdateCallback`:

C#

```
namespace Ice {
    public interface PropertiesAdminUpdateCallback {
        void updated(System.Collections.Generic.Dictionary<string, string>
changes);
    }
}
```

The `updated` method receives a `dictionary<string, string>` value representing the properties that were added, changed or removed, with removed properties denoted by an entry whose value is an empty string. It is legal for the `updated` implementation to modify the set of callbacks, and Ice ignores any exceptions it might raise.

The following code demonstrates how to register a callback:

Copyright 2016, ZeroC, Inc.

C#

```
Ice.Object obj = communicator.findAdminFacet("Properties");
if(obj != null) { // May be nil if the facet is not enabled
    Ice.NativePropertiesAdmin facet = obj as Ice.NativePropertiesAdmin;
    Ice.PropertiesAdminUpdateCallback myCallback = new ...;
    facet.addUpdateCallback(myCallback);
}
```

See Also

- Properties and Configuration
- Versioning
- IceGrid

Copyright 2016, ZeroC, Inc.

The Logger Facet

On this page:

- The RemoteLogger Interface
- The LoggerAdmin Interface
- Configuring the Logger Facet

The `RemoteLogger` Interface

An administrator may find it useful to view the log of a running Ice application, without going through an intermediary file. This is especially useful for Ice services that use a system log, such as the Windows Event Log.

The `Logger` facet allows remote applications (such as administrative clients) to attach one or more remote loggers to the local logger of any Ice application (provided this application enables the admin object with a `Logger` facet). The implementation of the `Logger` facet installs its own logger, which intercepts the log messages sent to the local logger, caches the most recent log messages, and forwards these log messages (after optional filtering) to the attached remote loggers.

A remote logger is an Ice object that implements the `Ice::RemoteLogger` interface. Such object is typically implemented by an administrative client.

Slice

```
module Ice {

    enum LogMessageType { PrintMessage, TraceMessage, WarningMessage, ErrorMessage
};
    sequence<LogMessageType> LogMessageTypeSeq;
    struct LogMessage {
       LogMessageType type;
       long timestamp;
       string traceCategory;
       string message;
    };
    sequence<LogMessage> LogMessageSeq;

    interface RemoteLogger {
        void init(string prefix, LogMessageSeq logMessages);
        void log(LogMessage message);
    };
};
```

The `LogMessage` `struct` represents log messages sent to a local logger. Its `timestamp` data member is the number of microseconds since the Unix Epoch (January 1st, 1970 at 0:00 UTC).

When a remote logger is attached to a local logger, its `init` operation is called with the local logger's prefix and a list of recent log messages (see The LoggerAdmin Interface below). Then, each time a log message is sent to the local logger, the Logger facet forwards this message to the remote logger's `log` operation.

The Logger facet does not guarantee that `init` will be called on a remote logger before `log` is called on this remote logger, even though the log messages sent to `init` are always older than the log messages sent to `log`. It is indeed common for `log` to be called several times before `init` in applications that generate many logs. An implementation of `RemoteLogger` needs to handle this situation correctly: it can for example keep all log messages received before `init` in a queue, and later append this queue to the log messages received through `init`.

The `LoggerAdmin` Interface

Copyright 2016, ZeroC, Inc.

The `Logger` facet implements the `Ice::LoggerAdmin` interface:

Slice

```
module Ice {
    interface LoggerAdmin {

        void attachRemoteLogger(RemoteLogger* prx, LogMessageTypeSeq messageTypes,
StringSeq traceCategories, int messageMax)
            throws RemoteLoggerAlreadyAttachedException;

        bool detachRemoteLogger(RemoteLogger* prx);

        LogMessageSeq getLog(LogMessageTypeSeq messageTypes, StringSeq
traceCategories, int messageMax, out string prefix);
};
```

The operation `attachRemoteLogger` attaches a remote logger (`prx`) to the local logger with the following optional filters:

- `messageTypes` specifies the types of log messages this remote logger wants to receive. An empty sequence means no filtering–the `Logger` facet will forward all log message types.
- `traceCategories` is a sequence of categories for trace messages. An empty sequence means no filtering–the `Logger` facet will forward all trace categories.
- `messageMax` is the maximum number of log messages sent to init. If `messageMax` is negative, all available log messages will be sent to `init` (provided they satisfy the `messageTypes` and `traceCategories` filters). If messageMax is 0, no log message will be sent to `init`. If `messageMax` is greater than 0, the most recent `messageMax` log messages that satisfy the `messageTypes` and `traceCategories` filters will be sent to `init`.

For example, you can attach a remote logger with no filtering at all as follows:

C++

```
Ice::RemoteLoggerPrx remoteLogger = // my remote logger, typically an object
implemented by the local application
Ice::ObjectPrx admin = // get proxy to the admin object for the target application
Ice::LoggerAdminPrx loggerAdmin = Ice::LoggerAdminPrx::uncheckedCast(admin,
"Logger");
loggerAdmin->attachRemoteLogger(remoteLogger, Ice::LogMessageTypeSeq(),
Ice::StringSeq(), -1);
```

If you are interested only in errors and trace messages for category `Network` with no more than 10 such messages sent to init, you could write instead:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::RemoteLoggerPrx remoteLogger = // my remote logger, typically an object
implemented by the local application
Ice::ObjectPrx admin = // get proxy to the admin object for the target application
Ice::LoggerAdminPrx loggerAdmin = Ice::LoggerAdminPrx::uncheckedCast(admin,
"Logger");

Ice::LogMessageTypeSeq messageTypes;
messageTypes.push_back(Ice::ErrorMessage);
messageTypes.push_back(Ice::TraceMessage);

Ice::StringSeq traceCategories;
traceCategories.push_back("Network");

loggerAdmin->attachRemoteLogger(remoteLogger, messageTypes, traceCategories, 10);
```

The operation `detachRemoteLogger` detaches a remote logger (`prx`) from the local logger. `prx` does not need to match exactly the proxy provided to `attachRemoteLogger`: the `Logger` facet uses only the identity in this proxy. `detachRemoteLogger` returns true if this remote logger was found (and is now detached), and false otherwise.

The Logger facet automatically detaches a remote logger when a request sent to this remote logger fails. As a result, if an administrative client forgets to call `detachRemoteLogger` after destroying its remote logger, or crashes, this remote logger will remain attached only as long as it's not used.

The operation `getLog` retrieves the most recent `messageMax` log messages that match the optional `messageTypes` and `traceCategories` filters. These parameters have the same meaning as in the `attachRemoteLogger` operation described above. For example, if you want to retrieve only the 10 most recent errors and trace messages for category `Network`, you could write:

C++

```
Ice::ObjectPrx admin = // get proxy to the admin object for the target application
Ice::LoggerAdminPrx loggerAdmin = Ice::LoggerAdminPrx::uncheckedCast(admin,
"Logger");

Ice::LogMessageTypeSeq messageTypes;
messageTypes.push_back(Ice::ErrorMessage);
messageTypes.push_back(Ice::TraceMessage);

Ice::StringSeq traceCategories;
traceCategories.push_back("Network");

string prefix;
Ice::LogMessageSeq logMessages = loggerAdmin->getLog(messageTypes, traceCategories,
10, prefix);
```

Configuring the Logger Facet

The `Logger` facet caches the most recent log messages sent to application's Logger, to be able to provide these log messages to remote loggers (in the `init` operation) and to administrative clients that call `getLog` (see the `LoggerAdmin` Interface above). You can configure

Copyright 2016, ZeroC, Inc.

how many log messages are cached by the Logger facet with the `Ice.Admin.Logger.KeepLogs` and `Ice.Admin.Logger.KeepTraces` properties. The default is to keep the most recent 100 log messages other than trace messages plus the most recent 100 trace messages.

See Also

- Logger Facility

Copyright 2016, ZeroC, Inc.

The Metrics Facet

The `Metrics` facet implements the Instrumentation Facility to provide convenient access to metrics for the Ice run time and select Ice services. The metrics provided by this facet include the number of threads currently running and their state, the number of connections, information on invocations and dispatch, as well as connection establishment and endpoint name resolution.

On this page:

- Metrics Terminology
- Metrics Types
- The MetricsAdmin Interface
- Obtaining the Local Metrics Facet
- Metrics Attributes

Metrics Terminology

- metric: an "analytical measurement intended to quantify the state of a system", recorded by the Ice run time, such as bytes sent over a connection.
- metric name: the name of the metric, such as "number of connections".
- metrics map: a collection of metrics objects.
- metrics view: a collection of metrics maps. A view contains metrics maps of different types (e.g., connections and threads). Several metrics views can be configured with different purposes. For example, you can have a "debug" metrics view to get detailed metrics of each of the instrumented objects in the Ice communicator. This view can be enabled from time to time for debugging purposes but it's disabled most of the time. You could also have a more coarse-grained metrics view to collect data at a higher level, such as the amount of bytes received and sent by all the connections from the communicator. This metrics view can be enabled all the time.

Metrics Types

Metrics are specified as Slice classes defined in the `Ice/Metrics.ice` Slice file. All the metrics types are defined in the `IceMX` module.

The base class is `IceMX::Metrics`:

Slice

```
class Metrics
{
    string id;
    long total = 0;
    int current = 0;
    long totalLifetime = 0;
    int failures = 0;
};
```

A metrics object is an instance of `IceMX::Metrics` and represents metrics of one or more instrumented objects. An instrumented object can be anything that supports instrumentation. The Ice run time supports instrumentation of the following objects and activities:

- Threads
- Connections
- Invocations
- Dispatch
- Connection establishment
- Endpoint resolution

The `id` of a metric identifies the instrumented object(s). The `total` and `current` members specify the total and current number of instrumented objects created since the creation of the Ice communicator, respectively. The `totalLifetime` member is the sum of the lifetime of each instrumented object and `failures` is the number of failures that have occurred for the metrics object(s).

Failures are specified using a separate `IceMX::MetricsFailures` structure:

Copyright 2016, ZeroC, Inc.

Slice

```
struct MetricsFailures
{
    string id;
    StringIntDict failures;
};
```

The `failures` dictionary provides the count of each type of failure for a metric identified by `id`. Failures are dependent on the instrumented objects. For example, failures for Ice connections are represented with the name of the exception that caused the connection to fail (e.g., `Ice::ConnectionLostException` or `Ice::TimeoutException`).

A metrics map is simply defined as a sequence of `IceMX::Metrics` objects, and a metrics view is defined as a dictionary of metrics map:

Slice

```
sequence<Metrics> MetricsMap;
dictionary<string, MetricsMap> MetricsView;
```

The key for the metrics view dictionary is a string that identifies the metrics map. The Ice run time supports the following metrics maps:

Metrics map name	Slice class	Description
Connection	`IceMX::ConnectionMetrics`	Connection metrics.
Thread	`IceMX::ThreadMetrics`	Thread metrics. The Ice run time instruments threads for the communicator's thread pools as well as threads used internally.
Invocation	`IceMX::InvocationMetrics`	Client-side proxy invocation metrics.
Dispatch	`IceMX::DispatchMetrics`	Server-side dispatch metrics.
EndpointLookup	`IceMX::Metrics`	Endpoint lookup metrics. For tcp, ssl and udp endpoints, this corresponds to the DNS lookups made to resolve the host names in endpoints.
ConnectionEstablishment	`IceMX::Metrics`	Connection establishment metrics.

A metrics map can also contain sub-metrics maps. An example is the `Invocation` metrics map, which provides a `Remote` sub-metrics map to record metrics associated with remote invocations. The Slice class for remote invocation metrics is `IceMX::Metrics`.

The `MetricsAdmin` Interface

The Slice interface `IceMX::MetricsAdmin` allows you to retrieve the metrics associated with the Ice communicator:

Copyright 2016, ZeroC, Inc.

Slice

```
module IceMX {
exception UnknownMetricsView {};

interface MetricsAdmin {
    Ice::StringSeq getMetricsViewNames(out Ice::StringSeq disableViews);
    void enableMetricsView(string name)
        throws UnknownMetricsView;
    void disableMetricsView(string name)
        throws UnknownMetricsView;
    MetricsView getMetricsView(string view, out long timestamp)
        throws UnknownMetricsView;
    MetricsFailuresSeq getMapMetricsFailures(string view, string map)
        throws UnknownMetricsView;
    MetricsFailures getMetricsFailures(string view, string map, string id)
        throws UnknownMetricsView;
};
};
```

The `getMetricsViewName` operation retrieves the names of the configured enabled and disabled views. The enableMetricsView and disableMetricsView allow to enable and disable a specific view. Calling those methods is equivalent to setting the view Disabled property to 1 or 0.The `getMetricsView` operation returns the metrics for the given view. The `getMapMetricsFailures` and `getMetricsFailures` operations retrieve the metrics failures for a given map or metrics id.

Obtaining the Local Metrics Facet

We already showed how to obtain a proxy for a remote administrative facet, but suppose you want to interact with the facet in your local address space. The code below shows the necessary steps:

C++

```
Ice::ObjectPtr obj = communicator->findAdminFacet("Metrics");
if(obj) { // May be nil if the facet is not enabled
    Ice::MetricsAdminPtr facet = Ice::MetricsAdminPtr::dynamicCast(obj);
    ...
}
```

As shown here, the facet is registered with the name `Metrics` and a regular language cast is used to downcast the base object type to the `MetricsAdmin` interface.

Metrics Attributes

Metrics views are configured with IceMX Metrics properties.

The `GroupBy`, `Accept` and `Reject` properties are specified using attributes that are specific to each metrics map. The table below describes the attributes supported by the Ice run time's metrics maps.

Name	Maps	Description
id	All	A unique identifier to select the instrumented object or operation.

Copyright 2016, ZeroC, Inc.

parent	All	The parent of the instrumented object or operation.
none	All	The none attribute is a special attribute that evaluates to the empty string.
endpoint	Connection, Dispatch, Remote, ConnectionEstablishment, EndpointLookup	The stringified endpoint.
endpointType	Connection, Dispatch, Remote, ConnectionEstablishment, EndpointLookup	The endpoint numerical type as defined in `Ice/Endpoint.ice`.
endpointIsDatagram	Connection, Dispatch, Remote, ConnectionEstablishment, EndpointLookup	A boolean indicating if the endpoint is a datagram endpoint.
endpointIsSecure	Connection, Dispatch, Remote, ConnectionEstablishment, EndpointLookup	A boolean indicating if the endpoint is secure.
endpointTimeout	Connection, Dispatch, Remote, ConnectionEstablishment, EndpointLookup	The endpoint timeout.
endpointCompress	Connection, Dispatch, Remote, ConnectionEstablishment, EndpointLookup	A boolean indicating if the endpoint requires compression.
endpointHost	Connection, Dispatch, Remote, ConnectionEstablishment, EndpointLookup	The endpoint host.
endpointPort	Connection, Dispatch, Remote, ConnectionEstablishment, EndpointLookup	The endpoint port.
connection	Dispatch	The connection description.
incoming	Connection, Dispatch, Remote	A boolean where true indicates an incoming (server) connection and false an outgoing (client) connection.
adapterName	Connection, Dispatch, Remote	If the connection is a server connection, adapterName returns the name of the adapter that created the connection, otherwise it is the empty string.
connectionId	Connection, Dispatch, Remote	The ID of the connection if one is set, otherwise it is the empty string.
localAddress	Connection, Dispatch, Remote	The connection's local address.
localPort	Connection, Dispatch, Remote	The connection's local port.
remoteAddress	Connection, Dispatch, Remote	The connection's remote address.
remotePort	Connection, Dispatch, Remote	The connection's remote port.
mcastAddress	Connection, Dispatch, Remote	The connection's multicast address.
mcastPort	Connection, Dispatch, Remote	The connection's multicast port.
state	Connection	The state of the connection.
operation	Dispatch, Invocation	The dispatched or invoked operation name.
identity	Dispatch, Invocation	The identity of the Ice object used for the dispatch or invocation.
facet	Dispatch, Invocation	The facet of the Ice object used for the dispatch or invocation.
mode	Dispatch, Invocation	The dispatch or invocation mode.
context.*key*	Dispatch, Invocation	The value of the dispatch or invocation context with the given key.
proxy	Invocation	The proxy used for the invocation.
encoding	Invocation	The proxy encoding.

Copyright 2016, ZeroC, Inc.

The `id`, `parent` and `none` attributes are supported by all maps.

The value of the `parent` attribute depends on the map. For the Connection, Dispatch and Remote maps, the `parent` will either be "Communicator" if the connection is a client connection, or the object adapter name if it's a server connection. For the Invocation, EndpointLookup, and ConnectionEstablishment maps, it will always be "Communicator". The `parent` attribute enables the filtering of metrics based on the object adapter. When used with the `GroupBy` property it also allows you to obtain metrics at the object adapter level. For instance, the following configuration does not monitor any metrics for the `Ice.Admin` object adapter and it groups all the metrics based on the object adapter or communicator:

```
IceMX.Metrics.MyView.GroupBy=parent
IceMX.Metrics.MyView.Reject.parent=Ice\.Admin   # Escape the dot in Ice.Admin
```

This configuration enables the communicator to get metrics on a per object adapter or communicator basis.

You can also use the `none` attribute to get metrics for the communicator including the metrics from object adapters, e.g., `IceMX.Metrics.MyView.GroupBy=none`. This provides the lowest possible level of detail as all the statistics will be recorded by a single metrics object.

The `id` attribute allows you to get a higher level of detail by recording metrics on a per instrumented object or operation basis. If you specify `IceMX.Metrics.MyView.GroupBy=id`, the `Metrics` facet will record metrics for each individual object or operation.

See Also

- Instrumentation Facility
- Creating the admin Object
- Ice.Admin.*

Copyright 2016, ZeroC, Inc.

Filtering Administrative Facets

The Ice run time enables all of its built-in administrative facets by default, and an application may install its own custom facets. You can control which facets the Ice run time enables using the `Ice.Admin.Facets` property. For example, the following property definition enables the `Properties` facet and leaves the `Process` facet (and any application-defined facets) disabled:

```
Ice.Admin.Facets=Properties
```

To specify more than one facet, separate them with a comma or white space. A facet whose name contains white space must be enclosed in single or double quotes.

The Ice run time creates only the built-in administrative facets that are enabled. Disabled built-in facets are not created at all. In Ice releases prior to Ice 3.6, the built-in facets were always created, whether enabled or disabled through `Ice.Admin.Facets`.

See Also

- The Process Facet
- The Properties Facet
- The Metrics Facet
- Custom Administrative Facets
- Ice.Admin.*

Copyright 2016, ZeroC, Inc.

Custom Administrative Facets

An application can add and remove administrative facets using the `Communicator` operations shown below:

Slice

```
module Ice {
local interface Communicator {
    // ...
    void addAdminFacet(Object servant, string facet);
    Object removeAdminFacet(string facet);
};
};
```

The `addAdminFacet` operation installs a new facet with the given name, or raises `AlreadyRegisteredException` if a facet already exists with the same name. The `removeAdminFacet` operation removes (and returns) the facet with the given name, or raises `NotRegisteredException` if no matching facet is found.

The mechanism for filtering administrative facets also applies to application-defined facets. If you call `addAdminFacet` while a filter is in effect, and the name of your custom facet does not match the filter, the Ice run time will not expose your facet but instead keeps a reference to it so that a subsequent call to `removeAdminFacet` is possible.

We provide an example of using these communicator operations in our discussion of the `Process` facet.

See Also

- Filtering Administrative Facets
- The Process Facet

Copyright 2016, ZeroC, Inc.

Security Considerations for Administrative Facets

Exposing administrative functionality naturally makes a program vulnerable, therefore it is important that proper precautions are taken. With respect to the default functionality, the `Properties` facet could expose sensitive configuration information, and the `Process` facet supports a `shutdown` operation that opens the door for a denial-of-service attack. Developers should carefully consider the security implications of any additional administrative facets that an application installs.

There are several approaches you can take to mitigate the possibility of abuse:

- Disable the administrative facility

  The administrative facility is disabled by default, and remains disabled as long as its prerequisites are not met. Note that IceGrid ena bles the facility in servers that it activates for the following reasons:
  - `Process`
  - `Properties`
  - `Logger`Loggers

  You could disable a facet using filtering, but doing so may disrupt IceGrid's normal operation.

- Select a proper endpoint

  If you host the admin object in the `Ice.Admin` object adapter, a reasonably secure value for the `Ice.Admin.Endpoints` property is one that uses the local host interface (`-h 127.0.0.1`), which restricts access to clients that run on the same host. Incidentally, this is the default value that IceGrid defines for its servers, although you can override that if you like. Note that using a local host endpoint does not preclude remote administration for IceGrid servers because IceGrid transparently routes requests on `admin` obje cts to the appropriate server via its node. If your application must support administration from non-local hosts, we recommend the use of SSL and certificate-based access control.

- Filter the facets

  After choosing a suitable endpoint, you can minimize risks by filtering the facets to enable only the functionality that is required. For example, if you are not using IceGrid's server activation feature and do not require the ability to remotely terminate a program, you should disable the `Process` facet using the filtering mechanism.

- Consider the object's identity

  The default identity of an `admin` object created during communicator initialization (or with `getAdmin`) has a UUID for its category, which makes it difficult for a hostile client to guess. Depending on your requirements, the use of a UUID may be an advantage or a disadvantage. For example, in a trusted environment, the use of a UUID may create additional work, such as the need to add an interface that an administrative client can use to obtain the identity or proxy of a remote `admin` object. An obscure identity might be more of a hindrance in this situation, and therefore specifying a static category via the `Ice.Admin.InstanceName` property is a reasonable alternative. In general, however, we recommend using the default behavior.

See Also

- The admin Object
- The Properties Facet
- The Process Facet
- Creating the admin Object
- IceGrid and the Administrative Facility
- Filtering Administrative Facets
- Ice.Admin.*
- IceGrid
- IceSSL

Copyright 2016, ZeroC, Inc.

Logger Facility

Depending on the setting of various properties, the Ice run time produces trace, warning, or error messages. These messages are written via the `Ice::Logger` interface:

Slice

```
module Ice {
    local interface Logger {
        void print(string message);
        void trace(string category, string message);
        void warning(string message);
        void error(string message);
        string getPrefix();
        Logger cloneWithPrefix(string prefix);
    };
};
```

The `cloneWithPrefix` operation returns a new logger that logs to the same destination but with a different prefix. (The prefix is used to, for example, provide the name of the process writing the log messages.)

Topics

- The Default Logger
- Custom Loggers
- Built-in Loggers
- Logger Plug-ins
- The Per-Process Logger
- C++ Logger Utility Classes

See Also

- Properties and Configuration

Copyright 2016, ZeroC, Inc.

The Default Logger

A default logger is instantiated when you create a communicator. The default logger writes its messages to the standard error output. The `trace` operation accepts a `category` parameter in addition to the error message; this allows you to separate trace output from different subsystems by sending the output through a filter.

You can obtain the logger that is attached to a communicator using the `getLogger` operation:

Slice

```
module Ice {
    local interface Communicator {
        Logger getLogger();
        // ...
    };
};
```

On Windows, when the default C++ logger outputs a log message to the console, it converts this message from your narrow string encoding (as defined by narrow string converter you installed, if any) to your console code page. This conversion is disabled if you redirect standard error to a file with `Ice.StdErr`, or if you set `Ice.LogStdErr.Convert` to 0.

See Also

- Logger Facility
- C++ Strings and Character Encoding

Copyright 2016, ZeroC, Inc.

Custom Loggers

You have several options if you wish to install a logger other than the default one:

- Select one of the other built-in loggers, which allow you to log to a file, to the `syslog` on Unix, and to the Windows event log
- Supply your own logger implementation in an `InitializationData` parameter when you create a communicator
- Load a logger implementation dynamically via the Ice plug-in facility

Changing the `Logger` object that is attached to a communicator allows you to integrate Ice messages into your own message handling system. For example, for a complex application, you might have an existing logging framework. To integrate Ice messages into that framework, you can create your own `Logger` implementation that logs messages to the existing framework.

When you destroy a communicator, its logger is *not* destroyed. This means that you can safely use a logger even beyond the lifetime of its communicator.

See Also

- Built-in Loggers
- Communicator Initialization
- Logger Plug-ins

Copyright 2016, ZeroC, Inc.

Built-in Loggers

Ice provides a file-based logger as well as Unix- and Windows-specific logger implementations. For .NET, the default Ice logger uses a `TraceListener` and so can be customized at run time via configuration.

On this page:

- File Logger
- Syslog Logger
- Windows Logger
- .NET Logger

File Logger

The file-based logger is enabled via the `Ice.LogFile` property. This logger is available for all supported languages and platforms.

Syslog Logger

You can activate a logger that logs via the Unix `syslog` implementation by setting the `Ice.UseSyslog` property. This logger is available in Ice for C++, Java, and C#, as well as for scripting languages based on Ice for C++.

Windows Logger

On Windows, subclasses of `Ice::Service` use the Windows application event log by default. The event log implementation is available for C++ applications.

.NET Logger

The default logger in Ice for .NET writes its messages using the `System.Diagnostics.Trace` facility. By default, the Ice run time registers a `ConsoleTraceListener` that writes to `stderr`. You can disable the logging of messages via this trace listener by setting the property `Ice.ConsoleListener` to zero.

You can change the trace listener for your application via the application's configuration file. For example:

```
<configuration>
  <system.diagnostics>
    <trace autoflush="false" indentsize="4">
      <listeners>
        <add name="myListener"
          type="System.Diagnostics.EventLogTraceListener"
          initializeData="TraceListenerLog" />
      </listeners>
    </trace>
  </system.diagnostics>
</configuration>
```

This configuration installs a trace listener that logs to the Windows event log using the source name `TraceListenerLog`.

The `EventLogTraceListener` creates a new event source if no match is found for the source name defined by `initializeData`, and creating a new event source requires that you run the application with administrative privileges. Alternatively, you can create the event source in advance using an administrative tool. For more information, search MSDN for "EventLog Component".

See Also

- The Server-Side main Function in C++

Copyright 2016, ZeroC, Inc.

Logger Plug-ins

Installing a custom logger using the Ice plug-in facility has several advantages. Because the logger plug-in is specified by a configuration property and loaded dynamically by the Ice run time, an application requires no code changes in order to utilize a custom logger implementation. Furthermore, a logger plug-in takes precedence over the per-process logger and the logger supplied in the `InitializationData` argument during communicator initialization, meaning you can use a logger plug-in to override the logger that an application installs by default.

On this page:

- Installing a C++ Logger Plug-in
- Installing a Java Logger Plug-in
- Installing a C# Logger Plug-in

Installing a C++ Logger Plug-in

To install a logger plug-in in C++, you must first define a subclass of `Ice::Logger`:

C++

```
class MyLoggerI : public Ice::Logger {
public:

    virtual void print(const std::string& message);
    virtual void trace(const std::string& category, const std::string& message);
    virtual void warning(const std::string& message);
    virtual void error(const std::string& message);
    virtual std::string getPrefix();
    virtual LoggerPtr void cloneWithPrefix(const string& prefix);

    // ...
};
```

Next, supply a factory function that installs your custom logger by returning an instance of `Ice::LoggerPlugin`:

C++

```
extern "C"
{

ICE_DECLSPEC_EXPORT Ice::Plugin*
createLogger(const Ice::CommunicatorPtr& communicator,
             const std::string& name,
             const Ice::StringSeq& args)
{
    Ice::LoggerPtr logger = new MyLoggerI;
    return new Ice::LoggerPlugin(communicator, logger);
}

}
```

The factory function can have any name you wish; we used `createLogger` in this example. Refer to the plug-in API for more information on plug-in factories.

Copyright 2016, ZeroC, Inc.

The definition of `LoggerPlugin` is shown below:

C++

```
namespace Ice {
class LoggerPlugin {
public:
    LoggerPlugin(const CommunicatorPtr&, const LoggerPtr&);

    virtual void initialize();
    virtual void destroy();
};
}
```

The constructor installs your logger into the given communicator. The `initialize` and `destroy` methods are empty, but you can subclass `LoggerPlugin` and override these methods if necessary.

Finally, define a configuration property that loads your plug-in into an application:

```
Ice.Plugin.MyLogger=mylogger:createLogger
```

The plug-in's name in this example is `MyLogger`; again, you can use any name you wish. The value of the property represents the plug-in's entry point, in which `mylogger` is the abbreviated form of its shared library or DLL, and `createLogger` is the name of the factory function.

If the configuration file containing this property is shared by programs in multiple implementation languages, you can use an alternate syntax that is loaded only by the Ice for C++ run time:

```
Ice.Plugin.MyLogger.cpp=mylogger:createLogger
```

Installing a Java Logger Plug-in

To install a logger plug-in in Java, you must first define a subclass of `Ice.Logger`:

Java

```
public class MyLoggerI implements Ice.Logger {

    public void print(String message) { ... }
    public void trace(String category, String message) { ... }
    public void warning(String message) { ... }
    public void error(String message) { ... }
    public String getPrefix() { ... }
    public Logger cloneWithPrefix(String prefix) { ... }

    // ...
}
```


Copyright 2016, ZeroC, Inc.

Next, define a factory class that installs your custom logger by returning an instance of `Ice.LoggerPlugin`:

Java

```
public class MyLoggerPluginFactoryI implements Ice.PluginFactory {
    public Ice.Plugin create(Ice.Communicator communicator,
                             String name, String[] args)
    {
        Ice.Logger logger = new MyLoggerI();
        return new Ice.LoggerPlugin(communicator, logger);
    }
}
```

The factory class can have any name you wish; in this example, we used `MyLoggerPluginFactoryI`. Refer to the plug-in API for more information on plug-in factories.

The definition of `LoggerPlugin` is shown below:

Java

```
package Ice;

public class LoggerPlugin implements Plugin {
    public LoggerPlugin(Communicator communicator, Logger logger)
    {
        // ...
    }

    public void initialize() { }

    public void destroy() { }
}
```

The constructor installs your logger into the given communicator. The `initialize` and `destroy` methods are empty, but you can subclass `LoggerPlugin` and override these methods if necessary.

Finally, define a configuration property that loads your plug-in into an application:

```
Ice.Plugin.MyLogger=MyLoggerPluginFactoryI
```

The plug-in's name in this example is `MyLogger`; again, you can use any name you wish. The value of the property is the name of the factory class.

If the configuration file containing this property is shared by programs in multiple implementation languages, you can use an alternate syntax that is loaded only by the Ice for Java run time:

```
Ice.Plugin.MyLogger.java=MyLoggerPluginFactoryI
```


Copyright 2016, ZeroC, Inc.

Installing a C# Logger Plug-in

To install a logger plug-in in .NET, you must first define a subclass of `Ice.Logger`:

C#

```
public class MyLoggerI : Ice.Logger {

    public void print(string message) { ... }
    public void trace(string category, string message) { ... }
    public void warning(string message) { ... }
    public void error(string message) { ... }
    public string getPrefix() { ... }
    public Logger cloneWithPrefix(string prefix) { ... }

    // ...
}
```

Next, define a factory class that installs your custom logger by returning an instance of `Ice.LoggerPlugin`:

C#

```
public class MyLoggerPluginFactoryI : Ice.PluginFactory {
    public Ice.Plugin create(Ice.Communicator communicator,
                             string name, string[] args)
    {
        Ice.Logger logger = new MyLoggerI();
        return new Ice.LoggerPlugin(communicator, logger);
    }
}
```

The factory class can have any name you wish; in this example, we used `MyLoggerPluginFactoryI`. Refer to the plug-in API for more information on plug-in factories. Typically the logger implementation and the factory are compiled into a single assembly.

The definition of `LoggerPlugin` is shown below:

Copyright 2016, ZeroC, Inc.

C#

```
namespace Ice {
public partial class LoggerPlugin : Plugin {
    public LoggerPlugin(Communicator communicator, Logger logger)
    {
        // ...
    }

    public void initialize() { }

    public void destroy() { }
}
}
```

The constructor installs your logger into the given communicator. The `initialize` and `destroy` methods are empty, but you can subclass `LoggerPlugin` and override these methods if necessary.

Finally, define a configuration property that loads your plug-in into an application:

```
Ice.Plugin.MyLogger=mylogger.dll:MyLoggerPluginFactoryI
```

The plug-in's name in this example is `MyLogger`; again, you can use any name you wish. The value of the property is the entry point for the factory, consisting of an assembly name followed by the name of the factory class.

If the configuration file containing this property is shared by programs in multiple implementation languages, you can use an alternate syntax that is loaded only by the Ice for .NET run time:

```
Ice.Plugin.MyLogger.clr=mylogger.dll:MyLoggerPluginFactoryI
```

See Also

- Custom Loggers
- The Per-Process Logger
- Communicator Initialization
- Plug-in API
- Ice.Plugin.*
- Ice.InitPlugins
- Ice.PluginLoadOrder

Copyright 2016, ZeroC, Inc.

The Per-Process Logger

Ice allows you to install a per-process custom logger. This logger is used by all communicators that do not have their own specific logger established at the time a communicator is created.

You can set a per-process logger in C++ by calling `Ice::setProcessLogger`, and you can retrieve the per-process logger by calling `Ice::getProcessLogger`:

C++

```
LoggerPtr getProcessLogger();
void setProcessLogger(const LoggerPtr&);
```

If you call `getProcessLogger` without having called `setProcessLogger` first, the Ice run time installs a default per-process logger. Note that if you call `setProcessLogger`, only communicators created after that point will use this per-process logger; communicators created earlier use the logger that was in effect at the time they were created. (This also means that you can call `setProcessLogger` multiple times; communicators created after that point will use whatever logger was established by the last call to `setProcessLogger`.)

`getProcessLogger` and `setProcessLogger` are language-specific APIs that are not defined in Slice. Therefore, for Java and C#, these methods appear in the `Ice.Util` class.

For applications that use the `Application` or `Service` convenience classes and do not explicitly configure a logger, these classes set a default per-process logger that uses the `Ice.ProgramName` property as a prefix for log messages. The `Application` class is described in the server-side language mapping chapters; more information on the `Service` class can be found in The `Ice::Service` Class.

See Also

- Custom Loggers
- Communicator Initialization
- The Server-Side main Function in C++

Copyright 2016, ZeroC, Inc.

C++ Logger Utility Classes

The Ice run time supplies a collection of utility classes that make use of the logger facility simpler and more convenient. Each of the `Logger` interface's four operations has a corresponding helper class:

C++

```
namespace Ice {
class Print {
public:
    Print(const LoggerPtr&);
    void flush();
    ...
};

class Trace {
public:
    Trace(const LoggerPtr&, const std::string&);
    void flush();
    ...
};

class Warning {
public:
    Warning(const LoggerPtr&);
    void flush();
    ...
};

class Error {
public:
    Error(const LoggerPtr&);
    void flush();
    ...
};
}
```

The only notable difference among these classes is the extra argument to the `Trace` constructor; this argument represents the trace category.

To use one of the helper classes in your application, you simply instantiate it and compose your message:

C++

```
if (errorCondition) {
    Error err(communicator->getLogger());
    err << "encountered error condition: " << errorCondition;
}
```

The Ice run time defines the necessary stream insertion operators so that you can treat an instance of a helper class as if it were a standard C++ output stream. When the helper object is destroyed, its destructor logs the message you have composed. If you want to log more than one message using the same helper object, invoke the `flush` method on the object to log what you have composed so far and reset the object for a new message.

Copyright 2016, ZeroC, Inc.

The helper classes also supply insertion operators to simplify the task of logging an exception. The operators accept instances of `std::exception` (from which all Ice exceptions derive) and log the string returned by the `what` method. On some platforms, you can also enable the configuration property `Ice.PrintStackTraces`, which causes the helper classes to log the stack trace of the exception in addition to the value of `what`.

See Also

- Logger Facility

Copyright 2016, ZeroC, Inc.

Instrumentation Facility

The Ice run time can be instrumented using observer interfaces that monitor many aspects of the run time's internal objects and activities, including connections, threads, servant dispatching, proxy invocations, endpoint lookups and connection establishment. We refer to these internal objects and activities as "instrumented objects" in the discussion below. The application is responsible for implementing the instrumentation observer interfaces. Note however that an implementation of these interfaces is provided by the Metrics Facet, so most applications do not need to implement them and can instead collect metrics through the `MetricsAdmin` facet. The definition of the instrumentation interfaces can be found in the `Ice/Instrumentation.ice` Slice file.

The Ice run time uses the `Ice::Instrumentation::CommunicatorObserver` interface to obtain observers for instrumented objects created by the run time:

Slice

```
module Ice {
    module Instrumentation {
        local interface CommunicatorObserver {
            Observer getConnectionEstablishmentObserver(Endpoint endpt, string
connector);
            Observer getEndpointLookupObserver(Endpoint endpt);
            ConnectionObserver getConnectionObserver(ConnectionInfo c, Endpoint e,
                                                     ConnectionState s,
ConnectionObserver o);
            ThreadObserver getThreadObserver(string parent, string id, ThreadState
s,
                                             ThreadObserver o);
            InvocationObserver getInvocationObserver(Object* prx, string operation,
Context ctx);
            DispatchObserver getDispatchObserver(Current c);
            void setObserverUpdater(ObserverUpdater updater);
        };
    };
    local interface Communicator {
        Ice::Instrumentation::CommunicatorObserver getObserver();
        // ...
    };
};
```

The Ice run time calls the appropriate `get...Observer` operation each time a new instrumented object is created. The implementation of these methods should return an observer, or nil if the implementation does not want to monitor the instrumented object. This observer is associated with the instrumented object and receives notifications of any changes to its attributes or state. All observer interfaces derive from the `Ice::Instrumentation::Observer` base interface:

Slice

```
local interface Observer {
    void attach();
    void detach();
    void failed(string exceptionName);
};
```

The `attach` operation is called upon association of the observer with the new instrumented object. The `detach` operation is called when the object is destroyed. The `failed` operation is called to report any failures that might occur during the lifetime of the instrumented object. Observer specializations provide additional operations for monitoring other attributes. For example, here is the `Ice::Instrumentation::`

Copyright 2016, ZeroC, Inc.

`ConnectionObserver` interface:

Slice

```
local interface ConnectionObserver extends Observer {
    void sentBytes(int num);
    void receivedBytes(int num);
};
```

The `sentBytes` and `receivedBytes` methods are called by the Ice connection when new bytes are received or sent over the connection.

Shown below is an implementation of the communicator and connection observer interfaces to record sent and received bytes on a per-connection basis. The observer dumps how many bytes were received and sent for the connection when it is detached:

Copyright 2016, ZeroC, Inc.

C++

```
class ConnectionObserverImpl : public Ice::Instrumentation::ConnectionObserver {
public:
    ConnectionObserverImpl(const Ice::ConnectionInfoPtr& connInfo) :
        info(connInfo), sentBytes(0), receivedBytes(0)
    {
    }

    void attach() {}

    void detach()
    {
        cerr << info->remoteHost << ":" << info->remotePort << ": sent bytes = "
             << sentBytes << ", received bytes = " << receivedBytes << endl;
    }

    void sentBytes(int num)
    {
        sentBytes += num;
    }

    void receivedBytes(int num)
    {
        receivedBytes += num;
    }

    void failed(const std::string&)
    {
    }

private:

    Ice::ConnectionInfoPtr info;
    long sentBytes;
    long receivedBytes;
};

class CommunicatorObserverImpl : public Ice::Instrumentation::CommunicatorObserver
{
public:

    Ice::Instrumentation::ConnectionObserverPtr
    getConnectionObserver(const Ice::ConnectionInfoPtr& c, const Ice::EndpointPtr&
e,
                          Ice::Instrumentation::ConnectionState s,
                          const Ice::Instrumentation::ConnectionObserverPtr&
previous)
    {
        return new ConnectionObserverImpl(c);
    }
};
```


Copyright 2016, ZeroC, Inc.

For brevity we have omitted the implementation of the other `get...Observer` methods; they all return 0 as we are only interested in instrumenting connections.

To register your implementation, you must pass it in an `InitializationData` parameter when you initialize a communicator:

C++

```
Ice::InitializationData id;
id.observer = new CommunicatorObserverImpl();
Ice::CommunicatorPtr ic = Ice::initialize(id);
```

You can install a `CommunicatorObserver` object on either the client or the server side (or both). Here is some example output produced by installing our `CommunicatorObserver` and `ConnectionObserver` object implementations in a simple server:

```
127.0.0.1:3487: sent bytes = 14, received bytes = 32
127.0.0.1:3487: sent bytes = 33, received bytes = 14
127.0.0.1:3490: sent bytes = 14, received bytes = 14
...
```

In addition to the operations for retrieving observers, the `CommunicatorObserver` interface also defines a `setObserverUpdater` operation that is called by the Ice run time on initialization to provide an updater object to the `CommunicatorObserver` implementation. This updater object can be used to "refresh" some of the created observers. The updater object provided by the Ice run time implements the following interface:

Slice

```
local interface ObserverUpdater {
    void updateConnectionObservers();
    void updateThreadObservers();
};
```

The `CommunicatorObserver` implementation can call these operations to update the observers associated with Ice connections or threads. When one of these operations is called, the Ice run time calls the matching `get...Observer` method on the `CommunicatorObserver` interface for each of the instrumented objects. For example, if you call `updateConnectionObservers`, your implementation of `getConnectionObserver` will be called again for each Ice connection in the communicator. The `previous` parameter to `getConnectionObserver` represents the observer that is currently associated with the connection.

This mechanism can be used to re-configure the observers associated with instrumented objects. For instance, the application might not wish to instrument connections all the time but only when needed. It can use the observer updater to enable or disable the instrumentation. Here is the example above modified to provide this functionality:

Copyright 2016, ZeroC, Inc.

C++

```
class CommunicatorObserverImpl : public Ice::Instrumentation::CommunicatorObserver,
                                 private IceUtil::Mutex {
public:

    Ice::Instrumentation::ConnectionObserverPtr
    getConnectionObserver(const Ice::ConnectionInfoPtr& c, const Ice::EndpointPtr&
e,
                          Ice::Instrumentation::ConnectionState s,
                          const Ice::Instrumentation::ConnectionObserverPtr&
previous)
    {
        Lock sync(*this);
        return enabled ? new ConnectionObserverImpl(c) : 0;
    }

    void setEnabled(bool enabled)
    {
        {
            Lock sync(*this);
            if(this->enabled == enabled)
                return;
            this->enabled = enabled;
        }
        updater->updateConnections();
    }

    void setObserverUpdater(const Ice::Instrumentation::ObserverUpdaterPtr&
updater)
    {
        this->updater = updater;
    }

    const Ice::Instrumentation::ObserverUpdaterPtr updater;
    bool enabled;
};
```

As you can see in the example above, special care needs to be taken with respect to synchronization. The Ice run time can call observers with Ice internal locks held to guarantee consistency of the information passed to the `get...Observer` methods. It is therefore important that the implementation of your observers performs quickly and does not create deadlocks. Your observers should not make remote invocations or call Ice APIs that require acquiring locks on instrumented objects.

See Also

- Communicator Initialization
- The Metrics Facet

Copyright 2016, ZeroC, Inc.

IceBox

IceBox is an easy-to-use framework for Ice application services. The Service Configurator pattern [1] is a useful technique for configuring services and centralizing their administration. In practical terms, this means services are developed as dynamically-loadable components that can be configured into a general purpose "super server" in whatever combinations are necessary. IceBox is an implementation of the Service Configurator pattern for Ice services.

A generic IceBox server replaces the typical monolithic Ice server you normally write. The IceBox server is configured via properties with the application-specific services it is responsible for loading and managing, and it can be administered remotely. There are several advantages in using this architecture:

- Services loaded by the same IceBox server can be configured to take advantage of Ice's collocation optimizations. For example, if one service is a client of another service, and those services reside in the same IceBox server, then invocations between them can be optimized.
- Composing an application consisting of various services is done by configuration, not by compiling and linking. This decouples the service from the server, allowing services to be combined or separated as needed.
- Multiple Java services can be active in a single instance of a Java Virtual Machine (JVM). This conserves operating system resources when compared to running several monolithic servers, each in its own JVM.
- Services implement an IceBox service interface, providing a common framework for developers and a centralized administrative facility.
- IceBox support is integrated into IceGrid, the server activation and deployment service.

IceBox offers a refreshing change of perspective: developers focus on writing services, not applications. The definition of an application changes as well; using IceBox, an application becomes a collection of discrete services whose composition is determined dynamically by configuration, rather than statically by the linker.

Topics

- Developing IceBox Services
- Configuring IceBox Services
- Starting the IceBox Server
- IceBox Administration

References

1. Jain, P., et al. 1997. Dynamically Configuring Communication Services with the Service Configurator Pattern. C++ Report 9 (6).

Copyright 2016, ZeroC, Inc.

Developing IceBox Services

On this page:

- The IceBox Service Interface
- IceBox Service Example in C++
  - C++ Service Entry Point
- IceBox Service Example in Java
- IceBox Service Example in C#
- IceBox Service Failures

The IceBox `Service` Interface

Writing an IceBox service requires implementing the IceBox `Service` interface:

Slice

```
module IceBox {
local interface Service {
    void start(string name, Ice::Communicator communicator, Ice::StringSeq args);
    void stop();
};
};
```

As you can see, a service needs to implement only two operations, `start` and `stop`. These operations are invoked by the server; `start` is called after the service is loaded, and `stop` is called when the IceBox server is shutting down.

The `start` operation is the service's opportunity to initialize itself; this typically includes creating an object adapter and servants. The `name` a nd `args` parameters supply information from the service's configuration, and the `communicator` parameter is an `Ice::Communicator` ob ject created by the server for use by the service. Depending on the service configuration, this communicator instance may be shared by other services in the same IceBox server, therefore care should be taken to ensure that items such as object adapters are given unique names.

The `stop` operation must reclaim any resources used by the service. Generally, a service deactivates its object adapter, and may also need to invoke `waitForDeactivate` on the object adapter in order to ensure that all pending requests have been completed before the clean up process can proceed. The server is responsible for destroying the communicator instance that was passed to `start`.

Whether the service's implementation of `stop` should explicitly destroy its object adapter depends on other factors. For example, the adapter should be destroyed if the service uses a shared communicator, especially if the service could eventually be restarted. In other circumstances, the service can allow its adapter to be destroyed as part of the communicator's destruction.

These interfaces are declared as `local` for a reason: they represent a contract between the server and the service, and are not intended to be used by remote clients. Any interaction the service has with remote clients is done via servants created by the service.

IceBox Service Example in C++

The example we present here is taken from the `IceBox/hello` sample program provided in the Ice distribution.

The class definition for our service is quite straightforward, but there are a few aspects worth mentioning:

Copyright 2016, ZeroC, Inc.

C++

```
#include <IceBox/IceBox.h>

#if defined(_WIN32)
#   define HELLO_API __declspec(dllexport)
#else
#   define HELLO_API /**/
#endif

class HELLO_API HelloServiceI : public IceBox::Service {
public:
    virtual void start(const std::string&,
                       const Ice::CommunicatorPtr&,
                       const Ice::StringSeq&);
    virtual void stop();

private:
    Ice::ObjectAdapterPtr _adapter;
};
```

First, we include the IceBox header file so that we can derive our implementation from `IceBox::Service`. Second, the preprocessor definitions are necessary because, on Windows, this service resides in a Dynamic Link Library (DLL), therefore we need to export the class so that the server can load it properly.

The member definitions are equally straightforward:

Copyright 2016, ZeroC, Inc.

C++

```
#include <Ice/Ice.h>
#include <HelloServiceI.h>
#include <HelloI.h>

using namespace std;

void
HelloServiceI::start(
    const string& name,
    const Ice::CommunicatorPtr& communicator,
    const Ice::StringSeq& args)
{
    _adapter = communicator->createObjectAdapter(name);
    Ice::ObjectPtr object = new HelloI(communicator);
    _adapter->add(object, communicator->stringToIdentity("hello"));
    _adapter->activate();
}

void
HelloServiceI::stop()
{
    _adapter->deactivate();
}
```

The `start` method creates an object adapter with the same name as the service, activates a single servant of type `HelloI` (not shown), and activates the object adapter. The `stop` method simply deactivates the object adapter.

C++ Service Entry Point

The last piece of the puzzle is the *entry point* function, which the IceBox server calls to obtain an instance of the service:

C++

```
extern "C" {
    HELLO_API IceBox::Service*
    create(Ice::CommunicatorPtr communicator)
    {
        return new HelloServiceI;
    }
}
```

In this example, the `create` function returns a new instance of the `Hello` service. The name of the function is not important, but it must have the signature shown above. In particular, the function must have C linkage, accept a single parameter of type `Ice::CommunicatorPtr`, and return a native pointer to `IceBox::Service`.

C linkage is required so that the IceBox server can locate this function in a dynamically-loaded library. The restrictions imposed on functions with C linkage prevent us from using the normal Ice calling conventions, which never return native pointers and always pass smart pointers by const reference. For example, such a function cannot return an object type (such as a smart pointer), which forces us to return a native pointer instead.

Copyright 2016, ZeroC, Inc.

Configuring IceBox Services provides more information on entry points and describes how to configure your service into an IceBox server.

IceBox Service Example in Java

As with the C++ example presented in the previous section, the complete source for the Java example can be found in the `IceBox/hello` directory of the Ice distribution. The class definition for our service looks as follows:

Java

```
public class HelloServiceI implements IceBox.Service
{
    public void
    start(String name, Ice.Communicator communicator, String[] args)
    {
        _adapter = communicator.createObjectAdapter(name);
        Ice.Object object = new HelloI(communicator);
        _adapter.add(object, communicator.stringToIdentity("hello"));
        _adapter.activate();
    }

    public void
    stop()
    {
        _adapter.deactivate();
    }

    private Ice.ObjectAdapter _adapter;
}
```

The `start` method creates an object adapter with the same name as the service, activates a single servant of type `HelloI` (not shown), and activates the object adapter. The `stop` method simply deactivates the object adapter.

The server requires a service implementation to have a default constructor. This is the *entry point* for a Java IceBox service; that is, the server dynamically loads the service implementation class and invokes the default constructor to obtain an instance of the service.

This example is a trivial service, and yours will likely be much more interesting, but this does demonstrate how easy it is to write an IceBox service. After compiling the service implementation class, it can be configured into an IceBox server as described in Configuring IceBox Services.

IceBox Service Example in C#

The complete source for the C# example can be found in the `IceBox/hello` directory of the Ice distribution. The class definition for our service looks as follows:

Copyright 2016, ZeroC, Inc.

C#

```
class HelloServiceI : IceBox.Service
{
    public void
    start(string name, Ice.Communicator communicator, string[] args)
    {
        _adapter = communicator.createObjectAdapter(name);
        _adapter.add(new HelloI(), communicator.stringToIdentity("hello"));
        _adapter.activate();
    }

    public void
    stop()
    {
        _adapter.deactivate();
    }

    private Ice.ObjectAdapter _adapter;
}
```

The `start` method creates an object adapter with the same name as the service, activates a single servant of type `HelloI` (not shown), and activates the object adapter. The `stop` method simply deactivates the object adapter.

The server requires a service implementation to have a default constructor. This is the entry point for a C# IceBox service; that is, the server dynamically loads the service implementation class from an assembly and invokes the default constructor to obtain an instance of the service.

This example is a trivial service, and yours will likely be much more interesting, but this does demonstrate how easy it is to write an IceBox service. After compiling the service implementation class, it can be configured into an IceBox server as described in Configuring IceBox Services.

IceBox Service Failures

An exception raised by a service's implementation of its entry point, `start`, or `stop` methods causes IceBox to log a message. An exception that occurs during server startup also results in server termination.

A service implementation can indicate a failure by raising `IceBox::FailureException`:

Slice

```
module IceBox {
local exception FailureException {
    string reason;
};
};
```

Note that, as a local exception, C++ users must instantiate `FailureException` with file and line number information:

Copyright 2016, ZeroC, Inc.

C++

```
throw IceBox::FailureException(__FILE__, __LINE__, "my error message");
```

See Also

- Configuring IceBox Services
- Starting the IceBox Server

Copyright 2016, ZeroC, Inc.

Configuring IceBox Services

On this page:

- Installing an IceBox Service
- IceBox Service Configuration in C++
- IceBox Service Configuration in Java
- IceBox Service Configuration in C#
- Load Order for IceBox Services
- Using a Shared Communicator
- Inheriting Properties from the IceBox Server
- Logging Considerations for IceBox Services

Installing an IceBox Service

A service is configured into an IceBox server using a single `IceBox.Service` property. This property serves several purposes: it defines the name of the service, it provides the server with the service entry point, and it defines properties and arguments for the service.

The format of the property is shown below:

```
IceBox.Service.name=entry_point [args]
```

The *name* component of the property key is the service name (IceStorm, in this example). This name is passed to the service's `start` operation, and must be unique among all services configured in the same IceBox server. It is possible, though rarely necessary, to load two or more instances of the same service under different names.

The first argument in the property value is the entry point specification. Any arguments following the entry point specification are examined. If an argument has the form `--name=value`, then it is interpreted as a property definition that appears in the property set of the communicator passed to the service `start` operation. These arguments are removed, and any remaining arguments are passed to the `start` operation in the `args` parameter.

IceBox Service Configuration in C++

For a C++ service, the entry point must have the form *library[,version]:symbol*, where *library* is the simple name of the service's shared library or DLL, and *symbol* is the name of the entry point function. A "simple name" is one without any platform-specific prefixes or extensions; the server adds appropriate decorations depending on the platform. The simple name may include a leading path, and the version is optional. If specified, the version is embedded in the library name.

As an example, here is how we could configure IceStorm, which is implemented as an IceBox service in C++:

```
IceBox.Service.IceStorm=IceStormService,36:createIceStorm
```

IceBox uses the information provided in the entry point specification to compose a library name. For the IceStorm example shown above, IceBox on Windows would compose the library name `IceStormService36.dll`. If IceBox is compiled with debug information, it appends a `d` to the library name, so the name becomes `IceStormService36d.dll` instead.

The exact name of the library that is loaded depends on the naming conventions of the platform IceBox executes on. For example, on an OS X machine, the library name is `libIceStormService36.dylib`.

If the simple name does not include a leading path, the shared library or DLL must reside in a directory that appears in `PATH` on Windows or the shared library search path (such as `LD_LIBRARY_PATH`) on POSIX systems.

The entry point function, *symbol*, must have the signature that we originally presented in our example:

Copyright 2016, ZeroC, Inc.

C++

```
extern "C" IceBox::Service* function(Ice::CommunicatorPtr);
```

The communicator instance passed to this function is the IceBox server's communicator and should only be used for administrative purposes. For example, the entry point function could use this communicator's logger to display log messages. For a service's normal operations, it must use the communicator that it receives as an argument to its `start` method.

Here is a sample configuration for our C++ service:

```
IceBox.Service.Hello=HelloService:create --Ice.Trace.Network=1 hello there
```

This configuration results in the creation of a service named `Hello`. The service is expected to reside in `HelloService.dll` on Windows or `libHelloService.so` on Linux, and the entry point function `create` is invoked to create an instance of the service. The argument `--Ice.Trace.Network=1` is converted into a property definition, and the arguments `hello` and `there` become the two elements in the `args` sequence parameter that is passed to the `start` method.

IceBox Service Configuration in Java

For a Java service, the entry point is typically the class name (including any package) of the service implementation class, but may also include a leading path to a class directory or JAR file. The class must define a public constructor.

To instantiate the service, the IceBox server first checks to see if the service defines a constructor taking an argument of type `Ice.Communicator`. If so, the service invokes this constructor and passes the server's communicator, which should only be used for administrative purposes. For example, the constructor could use this communicator's logger to display log messages. For a service's normal operations, it must use the communicator that it receives as an argument to its `start` method.

If the service does not define a constructor taking an `Ice.Communicator` argument, the server invokes the service's default constructor.

Here is a sample configuration for our Java example:

```
IceBox.Service.Hello=HelloServiceI --Ice.Trace.Network=1 hello there
```

This configuration results in the creation of a service named `Hello`. The service is expected to reside in the class `HelloServiceI`. The argument `--Ice.Trace.Network=1` is converted into a property definition, and the arguments `hello` and `there` become the two elements in the `args` sequence parameter that is passed to the `start` method.

IceBox Service Configuration in C#

The entry point of a .NET service has the form *`assembly:class`*. The *`assembly`* component can be a partially or fully qualified assembly name, such as `hello,Version=0.0.0.0,Culture=neutral`, or an assembly DLL name such as `hello.dll` that may optionally include a leading relative or absolute path name.

The locations that are searched for the assembly depend on how you define the *`assembly`* component:

Value for *`assembly`*	Example	Semantics

Copyright 2016, ZeroC, Inc.

Fully-qualified assembly name (strong-named assembly)	`hello,Version=...,Culture=neutral,publicKeyToken=...`	1. Checks assemblies that have already been loaded 2. Searches the Global Assembly Cache (GAC) 3. Searches the directory containing the `iceboxnet` executable
Partially-qualified assembly name	`hello`	1. Checks assemblies that have already been loaded 2. Searches the directory containing the `iceboxnet` executable
Relative path name	`services\MyService.dll`	Path name is relative to `iceboxnet`'s current working directory. Be sure to include the `.dll` extension in the path name.
Absolute path name	`C:\services\MyService.dll`	Assembly must reside at the specified path name. Be sure to include the `.dll` extension in the path name.

See MSDN for more information on how the CLR locates assemblies.

The *`class`* component is the complete class name of the service implementation class, which must define a public constructor.

To instantiate the service, the IceBox server first checks to see if the service defines a constructor taking an argument of type `Ice.Communicator`. If so, the service invokes this constructor and passes the server's communicator, which should only be used for administrative purposes. For example, the constructor could use this communicator's logger to display log messages. For a service's normal operations, it must use the communicator that it receives as an argument to its `start` method.

If the service does not define a constructor taking an `Ice.Communicator` argument, the server invokes the service's default constructor.

Here is a sample configuration for our C# example:

```
IceBox.Service.Hello=helloservice.dll:HelloServiceI --Ice.Trace.Network=1 hello
there
```

This configuration results in the creation of a service named `Hello`. The service is expected to reside in the assembly named `helloservice.dll`, implemented by the class `HelloServiceI`. The argument `--Ice.Trace.Network=1` is converted into a property definition, and the arguments `hello` and `there` become the two elements in the `args` sequence parameter that is passed to the `start` method.

Load Order for IceBox Services

By default, the server loads the configured services in an undefined order, meaning services in the same IceBox server should not depend on one another. If services must be loaded in a particular order, the `IceBox.LoadOrder` property can be used:

```
IceBox.LoadOrder=Service1,Service2
```

In this example, `Service1` is loaded first, followed by `Service2`. Any remaining services are loaded after `Service2`, in an undefined order. Each service mentioned in `IceBox.LoadOrder` must have a matching `IceBox.Service` property.

During shutdown, services are stopped in the reverse of the order in which they were loaded.

Copyright 2016, ZeroC, Inc.

Using a Shared Communicator

IceBox creates a separate communicator instance for each service by default in order to minimize the chances of accidental conflicts among services. You can optionally specify that certain services use a shared communicator instead by setting `IceBox.UseSharedCommunicator.`*`name`* properties in the server's configuration:

```
IceBox.Service.Hello=...
IceBox.Service.Printer=...
IceBox.UseSharedCommunicator.Hello=1
IceBox.UseSharedCommunicator.Printer=1
```

A common use case for sharing a communicator between two or more services is enabling the use of collocation optimizations for invocations among those services. This optimization is not possible with the default behavior that creates a new communicator for each service.

IceBox prepares the property set of this shared communicator as follows:

- If services inherit the server's properties, the property set initially contains the server's properties (excluding `Ice.Admin.Endpoints`), otherwise the property set starts out empty.
- For each service that uses the shared communicator:
  - Merge its properties into the shared property set, overwriting any existing properties with the same names
  - Remove any properties from the shared property set that the service explicitly clears. For example, the definition `Ice.Trace.Network=` clears any existing setting of `Ice.Trace.Network` in the shared property set.
  - Translate service-specific command-line settings into properties (e.g., `--Hello.Debug=1`)

Service properties are merged in the same order as the services are loaded. As a result, the final value of a property that is defined by multiple services depends on the order in which those services are loaded. Let's expand our example to demonstrate this behavior:

```
# File: server.cfg
IceBox.Service.Hello=... --Ice.Config=hello.cfg --Hello.Debug=1
IceBox.Service.Printer=... --Ice.Config=printer.cfg
IceBox.UseSharedCommunicator.Hello=1
IceBox.UseSharedCommunicator.Printer=1
IceBox.InheritProperties=1
IceBox.LoadOrder=Hello,Printer
Ice.Trace.Network=1

# File: hello.cfg
Ice.Trace.Network=2

# File: printer.cfg
Ice.Trace.Network=3
```

The two services `Hello` and `Printer` use the shared communicator, and all services inherit the server's properties. As IceBox prepares the shared property set, the initial value of `Ice.Trace.Network` is `1` as defined in the (inherited) server's configuration. The `IceBox.LoadOrder` property specifies that the `Hello` service should be loaded first, therefore the `Ice.Trace.Network` property momentarily has the value 2 but ultimately has the value 3 after the properties for the `Printer` service are merged.

If we change the value of `IceBox.LoadOrder` so that IceBox loads `Printer` first, the value for `Ice.Trace.Network` in the shared communicator will be 2 instead because the setting in `hello.cfg` overrides all previous values.

Inheriting Properties from the IceBox Server

Copyright 2016, ZeroC, Inc.

By default, a service does not inherit the IceBox server's configuration properties. For example, consider the following server configuration:

```
IceBox.Service.Weather=... --Ice.Config=svc.cfg
Ice.Trace.Network=1
```

The `Weather` service only receives the properties that are defined in its `IceBox.Service` property. In the example above, the service's communicator is initialized with the properties from the file `svc.cfg`.

If services need to inherit the IceBox server's configuration properties, define the `IceBox.InheritProperties` property in the IceBox server's configuration:

```
IceBox.Service.Weather=... --Ice.Config=svc.cfg
Ice.Trace.Network=1
IceBox.InheritProperties=1
```

All services inherit the server's properties when `IceBox.InheritProperties` is set to a non-zero value. The service inherits all the properties of the IceBox servers, with the exception of properties whose names start with `Ice.Admin`.

The properties of the shared communicator are also affected by this setting.

Logging Considerations for IceBox Services

In Ice 3.4 and earlier, the IceBox server configures a variation of its own logger in the communicator that it creates for each service (the only difference being a service-specific logging prefix), therefore the logging scheme you configure for the IceBox server must be appropriate for all of its services. The only way a service can configure a different logger is by using a logger plug-in.

As of Ice 3.5, the IceBox server only configures a logger for a service if that service has not already specified its own logger via the Ice.LogFile or Ice.UseSyslog properties.

See Also

- IceBox.*
- Developing IceBox Services
- IceStorm
- Collocated Invocation and Dispatch
- Miscellaneous Ice.* Properties

Copyright 2016, ZeroC, Inc.

Starting the IceBox Server

Incorporating everything we discussed previously, we can now configure and start IceBox servers.

On this page:

- Starting the C++ IceBox Server
- Starting the Java IceBox Server
- Starting the C# IceBox Server
- IceBox Server Failures

Starting the C++ IceBox Server

The configuration file for our example C++ service is shown below:

```
IceBox.Service.Hello=HelloService:create
Hello.Endpoints=tcp -p 10001
```

Notice that we define an endpoint for the object adapter created by the `Hello` service.

Assuming these properties reside in a configuration file named `config`, we can usually start the C++ IceBox server as follows:

```
$ icebox --Ice.Config=config
```

Additional command line options are supported, including those that allow the server to run as a Windows service or Unix daemon.

Windows Notes
Several versions of the `icebox` executable are included in the binary distribution for Windows. In the table below, *<installdir>* represents the directory in which you installed the Ice binary distribution.

Executable	Notes
<installdir>\bin\icebox	Loads 32-bit IceBox services compiled with Visual Studio 2013
<installdir>\bin\iceboxd	Loads 32-bit debug IceBox services compiled with Visual Studio 2013
<installdir>\bin\x64\icebox	Loads 64-bit IceBox services compiled with Visual Studio 2013
<installdir>\bin\x64\iceboxd	Loads 64-bit debug IceBox services compiled with Visual Studio 2013
<installdir>\bin\vc110\icebox	Loads 32-bit IceBox services compiled with Visual Studio 2012
<installdir>\bin\vc110\iceboxd	Loads 32-bit debug IceBox services compiled with Visual Studio 2012
<installdir>\bin\vc110\x64\icebox	Loads 64-bit IceBox services compiled with Visual Studio 2012
<installdir>\bin\vc110\x64\iceboxd	Loads 64-bit debug IceBox services compiled with Visual Studio 2012

Linux Notes
Several versions of the `icebox` executable are included in the binary distributions for Linux platforms:

Executable	Notes
icebox	Loads 64-bit IceBox services; only available on 64-bit architectures
icebox++11	Loads 64-bit IceBox services compiled in C++11 or C++14 mode; only available on 64-bit architectures

Copyright 2016, ZeroC, Inc.

`icebox32`	Loads 32-bit IceBox services
`icebox32++11`	Loads 32-bit IceBox services compiled in C++11 or C++14 mode

Starting the Java IceBox Server

Our Java configuration is nearly identical to the C++ version, except for the entry point specification:

```
IceBox.Service.Hello=HelloServiceI
Hello.Endpoints=tcp -p 10001
```

Notice that we define an endpoint for the object adapter created by the `Hello` service.

Assuming these properties reside in a configuration file named `config`, we can start the Java IceBox server as follows:

```
$ java IceBox.Server --Ice.Config=config
```

Starting the C# IceBox Server

The configuration file for our example C# service is shown below:

```
IceBox.Service.Hello=helloservice.dll:HelloService
Hello.Endpoints=tcp -p 10001
```

Notice that we define an endpoint for the object adapter created by the `Hello` service.

Assuming these properties reside in a configuration file named `config`, we can start the C# IceBox server as follows:

```
$ iceboxnet --Ice.Config=config
```

IceBox Server Failures

At startup, an IceBox server inspects its configuration for all properties having the prefix `IceBox.Service` and initializes each service. If initialization fails for a service, the IceBox server invokes the `stop` operation on any initialized services, reports an error, and terminates.

See Also

- The Server-Side main Function in C++
- IceBox.*

Copyright 2016, ZeroC, Inc.

IceBox Administration

An IceBox server internally creates an object called the service manager that is responsible for loading and initializing the configured services. You can optionally expose this object to remote clients, such as the IceBox and IceGrid administrative utilities, so that they can execute certain administrative tasks.

On this page:

- IceBox Administrative Slice Interfaces
  - The IceBox ServiceManager Interface
  - The IceBox ServiceObserver Interface
- Enabling the Service Manager
- IceBox Admin Facets
- IceBox Administrative Client Configuration
- IceBox Administrative Utility

IceBox Administrative Slice Interfaces

The Slice definitions shown below comprise the IceBox administrative interface:

Slice

```
module IceBox {
exception AlreadyStartedException {};
exception AlreadyStoppedException {};
exception NoSuchServiceException {};

interface ServiceObserver {
    void servicesStarted(Ice::StringSeq services);
    void servicesStopped(Ice::StringSeq services);
};

interface ServiceManager {
    idempotent Ice::SliceChecksumDict getSliceChecksums();
    void startService(string service)
        throws AlreadyStartedException, NoSuchServiceException;
    void stopService(string service)
        throws AlreadyStoppedException, NoSuchServiceException;
    void addObserver(ServiceObserver* observer)
    void shutdown();
};
};
```

The IceBox `ServiceManager` Interface

The `ServiceManager` interface provides access to the service manager object of an IceBox server. It defines the following operations:

- `getSliceChecksums`
  Returns a dictionary of checksums that allows a client to verify that it is using the same Slice definitions as the server.
- `startService`
  Starts a pre-configured service that is currently inactive. This operation cannot be used to add new services at run time, nor will it cause an inactive service's implementation to be reloaded. If no matching service is found, the operation raises `NoSuchServiceException`. If the service is already active, the operation raises `AlreadyStartedException`.
- `stopService`

Copyright 2016, ZeroC, Inc.

Stops an active service but does not unload its implementation. The operation raises `NoSuchServiceException` if no matching service is found, and `AlreadyStoppedException` if the service is stopped at the time `stopService` is invoked.

- `addObserver`
  Adds an observer that is called when IceBox services are started or stopped. The service manager ignores operations that supply a null proxy, or a proxy that has already been registered.
- `shutdown`
  Terminates the services and shuts down the IceBox server.

The IceBox `ServiceObserver` Interface

An administrative client that is interested in receiving callbacks when IceBox services are started or stopped must implement the `ServiceObserver` interface and register the callback object's proxy with the service manager using its `addObserver` operation. The `ServiceObserver` interface defines two operations:

- `servicesStarted`
  Invoked immediately upon registration to supply the current list of active services, and thereafter each time a service is started.
- `servicesStopped`
  Invoked whenever a service is stopped, and when the IceBox server is shutting down.

The IceBox server unregisters an observer if the invocation of either operation causes an exception.

Our discussion of IceGrid includes an example that demonstrates how to register a `ServiceObserver` callback with an IceBox server deployed with IceGrid.

Enabling the Service Manager

IceBox's administrative functionality is disabled by default. You can enable it using the Ice administrative facility by defining endpoints for the `Ice.Admin` object adapter with the property `Ice.Admin.Endpoints`.

The `Ice.Admin` object adapter is enabled automatically in an IceBox server that is deployed by IceGrid.

With the administrative facility enabled, IceBox registers an administrative facet with the name `IceBox.ServiceManager`. We discuss the identity of the `admin` object below.

Exposing the service manager makes an IceBox server vulnerable to denial-of-service attacks from malicious clients. Consequently, you should choose the endpoints and transports carefully.

IceBox Admin Facets

When you enable the service manager, IceBox adds it as a facet of the server's `admin` object. As a result, the identity of the service manager is the same as that of the `admin` object, and the name of its facet is `IceBox.ServiceManager`.

The identity of the `admin` object uses either a UUID or a statically-configured value for its category, and the value `admin` for its name. For example, consider the following property definitions:

```
Ice.Admin.Endpoints=tcp -h 127.0.0.1 -p 10001
Ice.Admin.InstanceName=IceBox
```

In this case, the identity of the `admin` object is `IceBox/admin`.

IceBox also creates in each service communicator (shared communicator and per-service communicator) all the built-in facets enabled on its main communicator, and adds all these facets, except the `Process` facet, to its admin object. These facets are named `IceBox.Service.`*`service-name`*`.`*`facet-name`*, where *service-name* corresponds to the service name (for example `Hello` or `IceStorm`), and *facet-name* is

Copyright 2016, ZeroC, Inc.

the name of the built-in facet (for example `Properties` or `Logger`).

You can instruct IceBox to skip the admin facets for a specific service by setting the property `Ice.Admin.Enabled` to a numeric value (typically 0) in the configuration for that service.

IceBox Administrative Client Configuration

A client requiring administrative access to the service manager must first obtain (or be able to construct) a proxy for the `admin` object. The default identity of the `admin` object uses a UUID for its category, which means the client cannot predict the identity and therefore will be unable to construct the proxy itself. If the IceBox server is deployed with IceGrid, the client can use the technique described in our discussion of IceGrid to access its `admin` object.

In the absence of IceGrid, the IceBox server should set the `Ice.Admin.InstanceName` property if remote administration is required. In so doing, the identity of the `admin` object becomes well-known, and a client can construct the proxy on its own. For example, let's assume that the IceBox server defines the following property:

```
Ice.Admin.InstanceName=IceBox
```

A client can define the proxy for the `admin` object in a configuration property as follows:

```
ServiceManager.Proxy=IceBox/admin -f IceBox.ServiceManager -h 127.0.0.1 -p 10001
```

The proxy option `-f IceBox.ServiceManager` specifies the name of the service manager's administrative facet.

IceBox Administrative Utility

IceBox includes C++ and Java implementations of an administrative utility. The utilities have the same usage:

```
Usage: iceboxadmin [options] [command...]
Options:
-h, --help           Show this message.
-v, --version        Display the Ice version.

Commands:
start SERVICE        Start a service.
stop SERVICE         Stop a service.
shutdown             Shutdown the server.
```

The C++ utility is named `iceboxadmin`, while the Java utility is represented by the class `IceBox.Admin`.

The `start` command is equivalent to invoking `startService` on the service manager interface. Its purpose is to start a pre-configured service; it cannot be used to add new services at run time. Note that this command does not cause the service's implementation to be reloaded.

Similarly, the `stop` command stops the requested service but does not cause the IceBox server to unload the service's implementation.

The `shutdown` command stops all active services and shuts down the IceBox server.

The C++ and Java utilities obtain the service manager's proxy from the property `IceBoxAdmin.ServiceManager.Proxy`, therefore this proxy must be defined in the program's configuration file or on the command line, and the proxy's contents of depend on the server's

Copyright 2016, ZeroC, Inc.

configuration. If the IceBox server is deployed with IceGrid, we recommend using the IceGrid administrative utilities instead, which provide equivalent commands for administering an IceBox server. Otherwise, the proxy should have the endpoints and identity configured for the server.

See Also

- Slice Checksums
- Administrative Facility
- The admin Object
- The Properties Facet
- icegridadmin Command Line Tool
- IceGrid and the Administrative Facility
- IceBox.*
- IceBoxAdmin
- Ice.Admin.*

Copyright 2016, ZeroC, Inc.

Ice Plugins

Topics

- IceDiscovery
- IceLocatorDiscovery
- IceSSL

Copyright 2016, ZeroC, Inc.

IceDiscovery

IceDiscovery provides a location service using UDP multicast that enables Ice applications to discover objects and object adapters.

On this page:

- IceDiscovery Overview
  - IceDiscovery Concepts
  - IceDiscovery vs. IceGrid
- Installing IceDiscovery
- Configuring IceDiscovery
  - IceDiscovery Property Overview
  - Configuring IceDiscovery in Clients
  - Configuring IceDiscovery in Servers
  - Configuring a Locator Proxy
- Using IceDiscovery
  - IceDiscovery Design Decisions
  - IceDiscovery Sample Programs

IceDiscovery Overview

IceDiscovery is an Ice plug-in that must be installed in all of the clients and servers in your application. Once installed, IceDiscovery enables a client to dynamically locate objects using indirect proxies, which avoids the need for the client to statically configure the endpoints of the objects it uses. In a server, IceDiscovery makes objects and object adapters available for discovery with minimal additional effort.

IceDiscovery Concepts

This section reviews some concepts that will help you as you learn more about IceDiscovery.

Indirect Proxies

Indirect proxies have two formats:

- `identityOnly`
  This format uses only the object's identity.
- `identity@adapterId`
  This format combines the object identity and an adapter identifier. This identifier can either refer to a specific object adapter or a replica group.

Notice that neither format includes any endpoints, such as `tcp -h somehost -p 10000`. The ability to resolve an indirect proxy using only symbolic information, much like a DNS lookup, helps to loosen the coupling between clients and servers.

The Ice core delegates the resolution of indirect proxies to a standardized locator facility. This architecture offers a significant advantage: Ice applications can change locator settings via external configuration without requiring any changes to the application code.

Plug-in Facility

In addition to the locator facility, Ice also defines a standard plug-in facility through which applications can modify the functionality of the Ice core or add new capabilities, again using only external configuration.

IceDiscovery combines these two facilities: it's installed as a standard Ice plug-in via configuration, and this plug-in installs a custom locator implementation that enables discovery using UDP multicast.

We'll describe later how to install IceDiscovery in your clients and servers.

Replication

The locator facility includes support for replicated object adapters. For example, if `Adapter1` and `Adapter2` both participate in a replicate group identified as `TheGroup`, then the indirect proxy `someObject@TheGroup` could resolve to either `someObject@Adapter1` or `someObject@Adapter2`. Although there are two adapters that host `someObject` at the implementation level, both object adapters incarnate the same logical object. The application is responsible for ensuring that any persistent state is properly synchronized between the servers that host replicated object adapters.

Requests and Replies

Copyright 2016, ZeroC, Inc.

In general terms, here's what happens when a client uses an indirect proxy for the first time:

- The client's IceDiscovery plug-in transparently broadcasts a "lookup" request via multicast
- The lookup request includes the identity of the target object, the object adapter identifier (if any), and a proxy that the target server's IceDiscovery plug-in can use to communicate directly with the client's IceDiscovery plug-in
- Every program that installs the plug-in and uses the same addressing information receives the client's multicast lookup request; only the server that hosts the target object sends a reply
- The client plug-in waits for a reply using a configurable timeout period; if it doesn't receive a reply, it tries again a configurable number of times before giving up
- The target server's IceDiscovery plug-in sends a reply including a proxy for the target object, which the client's plug-in receives and returns to the client's application code
- At this point, the client's application code uses the returned proxy's endpoints to communicate directly with the target object using whichever transports its object adapter supports

Only the initial lookup request occurs over multicast. The reply from the server plug-in to the client plug-in occurs using UDP unicast (by default). All subsequent communication between the client and the target object proceed directly without intervention by the IceDiscovery plug-in.

IceDiscovery Domains

Nothing prevents two unrelated IceDiscovery applications from using the same multicast address and port, which means a plug-in from Application A can receive lookup requests from a client in Application B for objects and object adapters that might coincidentally match those of Application A. To avoid this situation, the applications should configure unique *domain identifiers*. The client plug-in includes its domain identifier in each lookup request it sends so that a server plug-in can ignore any requests that don't match its own domain identifier.

IceDiscovery vs. IceGrid

IceDiscovery and IceGrid both provide a location service but it helps to understand their differences when deciding which one to use in an application. Use IceDiscovery when your application needs a lightweight, transient location service. IceGrid's location service is backed by a persistent database and represents just one of the features that IceGrid offers, along with remote administration, on-demand server activation, and many others. If you need a location service but aren't yet ready to dive into IceGrid, start out using IceDiscovery; migrating to IceGrid later won't be difficult.

We provide a plug-in similar to IceDiscovery called IceLocatorDiscovery that integrates with IceGrid.

Installing IceDiscovery

The IceDiscovery plug-in must be installed in every client that needs to locate objects and in all of the servers that host those objects.

You can use the `Ice.Plugin` property to load and install the plug-in at runtime; the property value depends on the language mapping you're using:

```
# C++, Python, PHP, Ruby
Ice.Plugin.IceDiscovery=IceDiscovery:createIceDiscovery

# Java
Ice.Plugin.IceDiscovery=IceDiscovery:IceDiscovery.PluginFactory

# .NET
Ice.Plugin.IceDiscovery=IceDiscovery:IceDiscovery.PluginFactory
```

If using C++, instead of dynamically loading the plug-in at runtime your application can explicitly link with and register the plug-in. To register the plug-in, you must call the `Ice::registerIceDiscovery(bool loadOnInitialize = true)` function before the communicator initialization. The `loadOnInitialize` parameter specifies if the plug-in is installed when the communicator is initialized. If set to `false`, you will need to enable the plugin by setting the `Ice.Plugin.IceDiscovery` property to `1`.

`Ice::registerIceDiscovery` is a simple helper function that calls `Ice::registerPluginFactory`.

Copyright 2016, ZeroC, Inc.

Refer to the next section for information on configuring the plug-in.

Configuring IceDiscovery

Applications configure the IceDiscovery plug-in using configuration properties; the plug-in does not provide a local API.

IceDiscovery Property Overview

The IceDiscovery plug-in supports a number of configuration properties, most of which affect the endpoints that the plug-in uses to communicate with its peers:

- Lookup endpoint
  This is the multicast endpoint on which all lookup queries are broadcast. It must use an IPv4 or IPv6 address in the multicast range with a fixed port.
- Reply endpoint
  This is the endpoint on which a client receives replies from servers. In general, this endpoint should not use a fixed port.

The plug-in uses sensible default values for all of its configuration properties, such that it's often unnecessary to define any of the plug-in's properties. However, it's still important to understand how the plug-in derives its endpoint information.

IceDiscovery creates several object adapters in each communicator in which it's installed, including the object adapters `IceDiscovery.Multicast` and `IceDiscovery.Reply`. These object adapters correspond to the Lookup and Reply endpoints mentioned above, respectively. You can configure the endpoints of these object adapters directly by defining the properties `IceDiscovery.Multicast.Endpoints` and `IceDiscovery.Reply.Endpoints`. If you don't define an endpoint for an object adapter, the plug-in computes it as follows:

- `IceDiscovery.Multicast.Endpoints=udp -h` *`address`* `-p` *`port`* `[--interface` *`interface`*`]`
- `IceDiscovery.Reply.Endpoints=udp [--interface` *`interface`*`]`

where

- *`address`* is the value of `IceDiscovery.Address` - defaults to `239.255.0.1` if IPv4 is enabled or `ff15::1` if IPv4 is disabled
- *`port`* is the value of `IceDiscovery.Port` - defaults to `4061`
- *`interface`* is the value of `IceDiscovery.Interface`

Consequently, if you don't define any of these properties, the plug-in uses the following endpoints by default (assuming IPv4):

- `IceDiscovery.Multicast.Endpoints=udp -h 239.255.0.1 -p 4061`
- `IceDiscovery.Reply.Endpoints=udp`

Finally, you can also override the default endpoint that a client uses to broadcast its lookup queries by defining `IceDiscovery.Lookup`, otherwise the plug-in computes this endpoint as follows:

- `IceDiscovery.Lookup=udp -h` *`address`* `-p` *`port`* `[--interface` *`interface`*`]`

This endpoint must use the same address and port as `IceDiscovery.Multicast.Endpoints`.

As you can see, the properties `IceDiscovery.Address`, `IceDiscovery.Port` and `IceDiscovery.Interface` are simply used as convenient shortcuts for customizing the details of the plug-in's endpoints. For example, suppose we want to use a different multicast address and port:

```
IceDiscovery.Address=239.255.0.99
IceDiscovery.Port=8000
```

The plug-in derives the following properties from these settings:

```
IceDiscovery.Multicast.Endpoints=udp -h 239.255.0.99 -p 8000
IceDiscovery.Lookup=udp -h 239.255.0.99 -p 8000
```

All of the clients and servers comprising an application must use the same values for `IceDiscovery.Address` and `IceDiscovery.Port`. You should also consider defining `IceDiscovery.DomainId` to avoid any potential collisions from unrelated

Copyright 2016, ZeroC, Inc.

applications that happen to use the same address and port.

Configuring IceDiscovery in Clients

Aside from installing the plug-in and optionally configuring its addressing information, no other configuration steps are required for an IceDiscovery client.

Configuring IceDiscovery in Servers

In addition to installing the plug-in and optionally configuring its addressing information, you also need to configure an identifier for each of a server's object adapters that hosts well-known (discoverable) objects. For example, suppose a server creates an object adapter named `Hello` and we want its objects to be discoverable. We can configure the object adapter's `AdapterId` property as follows:

```
Hello.AdapterId=HelloAdapter
```

The identifier you select must be globally unique among the servers sharing the same address and domain settings.

To use object adapter replication, you'll need to include the `ReplicaGroupId` property for each replicated object adapter:

```
Hello.AdapterId=Hello1
Hello.ReplicaGroupId=HelloAdapter
```

The indirect proxy `someObject@Hello1` refers to an object in this particular object adapter, whereas the indirect proxy `someObject@HelloAdapter` could refer to any object having the identity `someObject` in any of the object adapters participating in the replica group `HelloAdapter`.

Configuring a Locator Proxy

The IceDiscovery plug-in calls `setDefaultLocator` on its communicator at startup, therefore it's not necessary for you to configure a locator proxy.

Using IceDiscovery

As its name implies, the IceDiscovery plug-in enables clients to locate objects at run time, as long as the servers hosting those objects are actively running and using the same configuration settings for address, port, domain, and so on. Since IceDiscovery relies on UDP multicast to broadcast the lookup requests, you'll need to ensure that your network supports this transport.

IceDiscovery Design Decisions

From a design perspective, incorporating IceDiscovery into your application requires answering the following questions:

- What is the set of well-known objects that will be available for discovery?
  Applications typically don't need to make *every* object available for discovery. Rather, designs often only make "bootstrap" or "factory" objects available for discovery, while proxies for other objects can be obtained by invoking operations on these initial objects.
- How should clients refer to these well-known objects?
  As we explained earlier, proxies for well-known objects can take two forms: an identity by itself, or an identity with an adapter identifier, such as `factory` and `factory@AccountAdapter`, respectively. Clearly, using only identities places a greater burden on the application to ensure that they are globally unique among all of the clients and servers sharing the same address and domain settings. Including an object adapter identifier can help to further partition the object namespace, so that `factory@AccountAdapter` and `factory@AdminAdapter` represent distinct objects without requiring artificially unique identities such as `accountFactory` and `adminFactory`.
- Can unrelated applications share the same multicast address and port?
  If so, select unique domain identifiers for the applications and configure `IceDiscovery.DomainId` properties to avoid the potential for subtle bugs.
- Do you need replicated objects?

Copyright 2016, ZeroC, Inc.

Replicated object adapters can improve the fault tolerance of your application by allowing independent server processes to implement the same logical objects. Configure your servers as described above, and ensure your clients resolve indirect proxies using the replica group identifiers.

IceDiscovery provides enough flexibility to support a wide variety of application architectures. If you need additional functionality, consider using IceGrid instead. Note also that IceGrid supports its own version of IceDiscovery, so that migrating an existing IceDiscovery application should be straightforward.

IceDiscovery Sample Programs

The Ice distribution includes two IceDiscovery sample programs:

- `hello` - A basic client/server application
- `replication` - An application that demonstrates the use of object adapter replication

Refer to the `README` files in the demo source directories for more information on these examples.

See Also

- IceDiscovery.*
- Plug-in Facility
- Locators
- Well-Known Objects
- IceGrid

Copyright 2016, ZeroC, Inc.

IceLocatorDiscovery

This page describes IceLocatorDiscovery, an Ice plug-in that enables the discovery of IceGrid and custom locators via UDP multicast.

On this page:

- IceLocatorDiscovery Overview
- Installing IceLocatorDiscovery
- Configuring IceLocatorDiscovery
  - IceLocatorDiscovery Property Overview
  - Configuring IceLocatorDiscovery in User Applications
  - Configuring IceLocatorDiscovery in IceGrid Administrative Clients
  - Configuring IceLocatorDiscovery in an IceGrid Registry
  - Configuring IceLocatorDiscovery in an IceGrid Node

IceLocatorDiscovery Overview

IceLocatorDiscovery is an Ice plug-in introduced in Ice 3.6 that discovers IceGrid and custom locators on a network using UDP multicast. Once installed, the plug-in automatically and transparently issues a multicast query in an attempt to find one or more locators, collects the responses, and configures the Ice run time accordingly. The primary advantage of using IceLocatorDiscovery is that it eliminates the need to manually configure and maintain the `Ice.Default.Locator` property that has traditionally been necessary in prior Ice releases. It's even more helpful in a replicated IceGrid deployment consisting of a master replica and one or more slave replicas, where the `Ice.Default.Locator` property would normally include endpoints for some or all of the replicas. Avoiding the need to statically configure the locator endpoints relieves some of the administrative burden, simplifies deployment and configuration tasks, and adds more flexibility to your application designs.

You can think of IceLocatorDiscovery as an application-specific version of IceDiscovery geared primarily toward IceGrid users.

Installing IceLocatorDiscovery

The IceLocatorDiscovery plug-in must be installed in every client that needs to locate objects; you can optionally install it in IceGrid nodes and registry replicas.

You can use the `Ice.Plugin` property to install the plug-in; the property value depends on the language mapping you're using:

```
# C++, Python, PHP, Ruby
Ice.Plugin.IceLocatorDiscovery=IceLocatorDiscovery:createIceLocatorDiscovery

# Java
Ice.Plugin.IceLocatorDiscovery=IceLocatorDiscovery:IceLocatorDiscovery.PluginFactor
y

# .NET
Ice.Plugin.IceLocatorDiscovery=IceLocatorDiscovery:IceLocatorDiscovery.PluginFactor
y
```

If using C++, instead of dynamically loading the plug-in at runtime your application can explicitly link with and register the plug-in. To register the plug-in, you must call the `Ice::registerIceLocatorDiscovery(bool loadOnInitialize = true)` function before the communicator initialization. The `loadOnInitialize` parameter specifies if the plug-in is installed when the communicator is initialized. If set to `false`, you will need to enable the plugin by setting the `Ice.Plugin.IceLocatorDiscovery` property to `1`.

`Ice::registerIceLocatorDiscovery` is a simple helper function that calls `Ice::registerPluginFactory`.

Refer to the next section for information on configuring the plug-in.

Copyright 2016, ZeroC, Inc.

Configuring IceLocatorDiscovery

Applications configure the IceLocatorDiscovery plug-in using configuration properties; the plug-in does not provide a local API.

IceLocatorDiscovery Property Overview

The IceDiscovery plug-in supports a number of configuration properties, many of which affect the endpoints that the plug-in uses for its queries:

- Lookup endpoint
  This is the multicast endpoint on which all lookup queries are broadcast. It must use an IPv4 or IPv6 address in the multicast range with a fixed port.
- Reply endpoint
  This is the endpoint on which the plug-in receives replies from locators (an IceGrid registry is the most common example of a locator). In general, this endpoint should not use a fixed port.

The plug-in uses sensible default values for all of its configuration properties, such that it's often unnecessary to define any of the plug-in's properties. However, it's still important to understand how the plug-in derives its endpoint information.

First, you can override the default endpoint that the plug-in uses to broadcast its queries by defining `IceLocatorDiscovery.Lookup`, otherwise the plug-in computes this endpoint as follows:

```
IceLocatorDiscovery.Lookup=udp -h address -p port [--interface interface]
```

where

- *`address`* is the value of `IceLocatorDiscovery.Address` - defaults to `239.255.0.1` if IPv4 is enabled or `ff15::1` if IPv4 is disabled
- *`port`* is the value of `IceLocatorDiscovery.Port` - defaults to `4061`
- *`interface`* is the value of `IceLocatorDiscovery.Interface`

For IceGrid users, the lookup endpoint must use the same multicast address and port as `IceGrid.Registry.Discovery.Endpoints` in the registry configuration.

IceLocatorDiscovery also creates object adapters in each communicator in which it's installed, including the object adapter `IceLocatorDiscovery.Reply`. This object adapter corresponds to the Reply endpoint mentioned above.

As you can see, the properties `IceLocatorDiscovery.Address`, `IceLocatorDiscovery.Port` and `IceLocatorDiscovery.Interface` are simply used as convenient shortcuts for customizing the details of the plug-in's endpoints. For example, suppose we want to use a different multicast address and port:

```
IceLocatorDiscovery.Address=239.255.0.99
IceLocatorDiscovery.Port=8000
```

The plug-in derives the following property from these settings:

```
IceLocatorDiscovery.Lookup=udp -h 239.255.0.99 -p 8000
```

All of the components of an IceGrid application must use the same multicast address and port. You should also consider defining `IceLocatorDiscovery.InstanceName` to avoid any potential collisions from unrelated IceGrid applications that happen to use the same address and port.

Configuring IceLocatorDiscovery in User Applications

For a client application, remove any existing definition of `Ice.Default.Locator`, then install the plug-in and optionally configuring its addressing information.

For a server deployed with IceGrid, you normally don't need to install the IceLocatorDiscovery plug-in.

Copyright 2016, ZeroC, Inc.

Configuring IceLocatorDiscovery in IceGrid Administrative Clients

Support for multicast discovery is built into the command-line and graphical IceGrid administrative utilities, therefore you don't need to install the plug-in. Both utilities support configuration properties similar to the ones we described above for defining the multicast address and port.

Configuring IceLocatorDiscovery in an IceGrid Registry

An IceGrid registry does not need the plug-in if it's the master replica or the only registry in a deployment, although there's no harm in installing it. The plug-in is useful for slave replicas because it allows them to locate the current master without statically configuring the master's endpoints. Configure the plug-in for slave replicas just like you would for any C++ client.

IceGrid registries listen for multicast discovery queries by default, but you can disable this feature by setting `IceGrid.Registry.Discovery.Enabled` to zero.

If you've changed the default multicast address or port for IceLocatorDiscovery, you must also make corresponding changes to the configuration of each registry. The registry supports properties similar to those of IceLocatorDiscovery:

- `IceGrid.Registry.Discovery.Address`
- `IceGrid.Registry.Discovery.Port`
- `IceGrid.Registry.Discovery.Interface`

These properties influence the endpoint on which the registry listens for multicast discovery queries. If you don't override the endpoint by setting `IceGrid.Registry.Discovery.Endpoints`, the registry uses these properties to compute its endpoint as follows:

```
IceGrid.Registry.Discovery.Endpoints=udp -h address -p port [--interface interface]
```

You don't need to define any `IceGrid.Registry.Discovery.*` properties if you want the registry to listen for discovery queries on its default multicast address and port.

Configuring IceLocatorDiscovery in an IceGrid Node

The plug-in is useful for IceGrid nodes because it allows them to locate the registry without statically configuring the registry's endpoints. Configure nodes just like you would for any C++ client.

See Also

- IceLocatorDiscovery Properties
- Registry Replication
- Plug-in Facility

Copyright 2016, ZeroC, Inc.

IceSSL

Security is an important consideration for many distributed applications, both within corporate intranets as well as over untrusted networks, such as the Internet. The ability to protect sensitive information, ensure its integrity, and verify the identities of the communicating parties is essential for developing secure applications. With those goals in mind, Ice includes the IceSSL *plug-in* that provides these capabilities using the Secure Socket Layer (SSL) protocol.

Although security is an optional component of Ice, it is not an afterthought. The IceSSL plug-in integrates easily into existing Ice applications, in most cases requiring nothing more than configuration changes. Naturally, some additional effort is required to create the necessary security infrastructure for an application, but in many enterprises this work will have already been done.

IceSSL is available for C++, Java and .NET applications. Python, PHP and Ruby applications can use IceSSL for C++ via configuration.

On this page:

- Overview of SSL
- Public Key Infrastructure
- Requirements

Overview of SSL

The Secure Socket Layer (SSL) protocol is the de facto standard for secure network communication. Its support for authentication, non-repudiation, data integrity, and strong encryption makes it the logical choice for securing Ice applications.

SSL is the protocol [1] that enables Web browsers to conduct secure transactions and therefore is one of the most commonly used protocols for secure network communication. You do not need to know the technical details of the SSL protocol in order to use IceSSL successfully (and those details are outside the scope of this text). However, it would be helpful to have a high-level understanding of how the protocol works and the infrastructure required to support it.

SSL provides a secure environment for communication (without sacrificing too much performance) by combining a number of cryptographic techniques:

- public key encryption
- symmetric (shared key) encryption
- message authentication codes
- digital certificates

When a client establishes an SSL connection to a server, a *handshake* is performed. During a typical handshake, digital certificates that identify the communicating parties are validated, and symmetric keys are exchanged for encrypting the session traffic. Public key encryption, which is too slow to be used for the bulk of a session's data transfer, is used heavily during the handshaking phase. Once the handshake is complete, SSL uses message authentication codes to ensure data integrity, allowing the client and server to communicate at will with reasonable assurance that their messages are secure.

Public Key Infrastructure

Security requires trust, and public key cryptography by itself does nothing to establish trust. SSL addresses the issue of trust using Public Key Infrastructure (PKI) [2], which binds public keys to identities using certificates. A certificate *issuer* creates a certificate for an entity, called the *subject*. The subject is often a person, but it may also be a computer or a specific application. The subject's identity is represented by a *distinguished name*, which includes information such as the subject's name, organization and location. A certificate alone is not sufficient to establish the subject's identity, however, as anyone can create a certificate for a particular distinguished name.

In order to authenticate a certificate, we need a third-party to guarantee that the certificate belongs to the subject described by the distinguished name. This third party, called a Certificate Authority (CA), expresses this guarantee by using its own private key to sign the subject's certificate. Combining the CA's certificate with the subject's certificate forms a *certificate chain* that provides SSL with most of the information it needs to authenticate the remote peer. In many cases, the chain contains only the aforementioned two certificates, but it is also possible for the chain to be longer when the *root* CA issues a certificate that the subject may use to sign other certificates. Regardless of the length of the chain, this scheme can only work if we trust that the root CA has sufficiently verified the identity of the subject before issuing the certificate.

An implementation of the SSL protocol also needs to know which root CAs we trust. An application supplies that information as a list of certificates representing the trusted root CAs. With that list in hand, the SSL implementation authenticates a peer by obtaining the peer's certificate chain and examining it carefully for validity. If we view the chain as a hierarchy with the root CA certificate at the top and the peer's certificate at the bottom, we can describe SSL's validation activities as follows:

- The root CA certificate must be self-signed and be present among the application's trusted CA certificates.

Copyright 2016, ZeroC, Inc.

- All other certificates in the chain must be signed by the one immediately preceding it.
- The certificates must not be expired or revoked.

These tests certify that the chain is valid, but applications often require the chain to undergo more intensive scrutiny to determine whether the chain is trustworthy.

Commercial CAs exist to supply organizations with a reliable source of certificates, but in many cases a private CA is completely sufficient. You can create and manage your CA using freely-available tools, and in fact Ice includes a collection of utilities that simplify this process.

Depending on your implementation language, it may also possible to avoid the use of certificates altogether; encryption is still used to obscure the session traffic, but the benefits of authentication are sacrificed in favor of reduced complexity and administration.

Requirements

Integrating IceSSL into your application often requires no changes to your source code, but does involve the following administrative tasks:

- creating a public key infrastructure (if necessary)
- configuring the IceSSL plug-in
- modifying your application's configuration to install the IceSSL plug-in and use secure connections

The remainder of this discussion concentrates on plug-in configuration and programming.

Topics

- Using IceSSL
- Configuring IceSSL
- Programming IceSSL
- Advanced IceSSL Topics
- Setting up a Certificate Authority

References

1. Viega, J., et al. 2002. *Network Security with OpenSSL*. Sebastopol, CA: O'Reilly.
2. Housley, R., and T. Polk. 2001. *Planning for PKI: Best Practices Guide for Deploying Public Key Infrastructure*. Hoboken, NJ: Wiley.

Copyright 2016, ZeroC, Inc.

Using IceSSL

Incorporating IceSSL into your application requires installing the plug-in, configuring it according to your security requirements, and creating SSL endpoints.

On this page:

- Installing IceSSL
  - C++ Applications
  - Java Applications
  - .NET Applications
  - Ice Touch Applications
- Creating SSL Endpoints
- Endpoint Security Considerations

Installing IceSSL

Ice supports a generic plug-in facility that allows extensions (such as IceSSL) to be installed dynamically without changing the application source code. The `Ice.Plugin` property provides language-specific information that enables the Ice run time to install a plug-in.

C++ Applications

The executable code for the IceSSL C++ plug-in resides in a shared library on Unix and a dynamic link library (DLL) on Windows. The format for the `Ice.Plugin` property is shown below:

```
Ice.Plugin.IceSSL=IceSSL:createIceSSL
```

The last component of the property name (`IceSSL`) becomes the plug-in's official identifier for configuration purposes, but the IceSSL plug-in requires its identifier to be `IceSSL`. The property value `IceSSL:createIceSSL` is sufficient to allow the Ice run time to locate the IceSSL library (on both Unix and Windows) and initialize the plug-in. The only requirement is that the library reside in a directory that appears in the shared library path (`LD_LIBRARY_PATH` on most Unix platforms, `PATH` on Windows).

Instead of dynamically loading the plug-in at runtime your application can also explicitly link with and register the plug-in. To register the plug-in, you must call the `Ice::registerIceSSL(bool loadOnInitialize = true)` function before the communicator initialization. The `loadOnInitialize` parameter specifies if the plug-in is installed when the communicator is initialized. If set to `false`, you will need to enable the plugin by setting the `Ice.Plugin.IceSSL` property to `1`.

Additional configuration properties are usually necessary as well.

`Ice::registerIceSSL` is a simple helper function that calls `Ice::registerPluginFactory`.

Java Applications

The format for the `Ice.Plugin` property is shown below:

```
Ice.Plugin.IceSSL=IceSSL.PluginFactory
```

The last component of the property name (`IceSSL`) becomes the plug-in's official identifier for configuration purposes, but the IceSSL plug-in requires its identifier to be `IceSSL`. The property value `IceSSL.PluginFactory` is the name of a class that allows the Ice run time to initialize the plug-in. The IceSSL classes are included in `Ice.jar`, therefore no other changes to your `CLASSPATH` are necessary.

Additional configuration properties are usually necessary as well.

.NET Applications

The format for the `Ice.Plugin` property is shown below:

Copyright 2016, ZeroC, Inc.

```
Ice.Plugin.IceSSL=C:/Ice/bin/IceSSL.dll:IceSSL.PluginFactory
```

The last component of the property name (`IceSSL`) becomes the plug-in's official identifier for configuration purposes, but the IceSSL plug-in requires its identifier to be `IceSSL`. The property value contains the file name of the IceSSL assembly as well as the name of a class, `IceSSL.PluginFactory`, that allows the Ice run time to initialize the plug-in.

You may also specify a partially or fully qualified assembly name instead of the file name in an `Ice.Plugin` property. For example, you can use the following configuration to load the plug-in from the release version of the assembly provided in ZeroC's binary distribution:

```
Ice.Plugin.IceSSL=IceSSL,Version=3.6.0.0,Culture=neutral,PublicKeyToken=cdd571ade22
f2f16:IceSSL.PluginFactory
```

Alternatively, if you have built the Ice source code yourself, you can load the development version of the assembly as shown below:

```
Ice.Plugin.IceSSL=IceSSL,Version=3.6.0.0,Culture=neutral,PublicKeyToken=1f998c50fec
78381:IceSSL.PluginFactory
```

You *must* use a fully qualified assembly name to load the IceSSL plug-in from the Global Assembly Cache.

Additional configuration properties are usually necessary as well.

Ice Touch Applications

The IceSSL plug-in is included in the Ice Touch run time and installed automatically, therefore it is not necessary to explicitly load it. However, configuration properties are usually necessary.

Creating SSL Endpoints

Installing the IceSSL plug-in enables you to use a new protocol, `ssl`, in your endpoints. For example, the following adapter endpoint list creates a TCP endpoint, an SSL endpoint, and a UDP endpoint:

```
MyAdapter.Endpoints=tcp -p 4063:ssl -p 4064:udp -p 4063
```

As this example demonstrates, it is possible for a UDP endpoint to use the same port number as a TCP or SSL endpoint, because UDP is a different protocol and therefore has an indepdendent set of ports. It is not possible for a TCP endpoint and an SSL endpoint to use the same port number, because SSL is essentially a layer over TCP.

Using SSL in proxy endpoints is equally straightforward:

```
MyProxy=MyObject:tcp -p 4063:ssl -p 4064:udp -p 4063
```

Endpoint Security Considerations

Defining an object adapter's endpoints to use multiple protocols, as shown above, has obvious security implications. If your intent is to use SSL to protect session traffic and/or restrict access to the server, then you should only define SSL endpoints.

There can be situations, however, in which insecure endpoint protocols are advantageous. The figure below illustrates an environment in which TCP endpoints are allowed behind the firewall, but external clients are required to use SSL:

Copyright 2016, ZeroC, Inc.

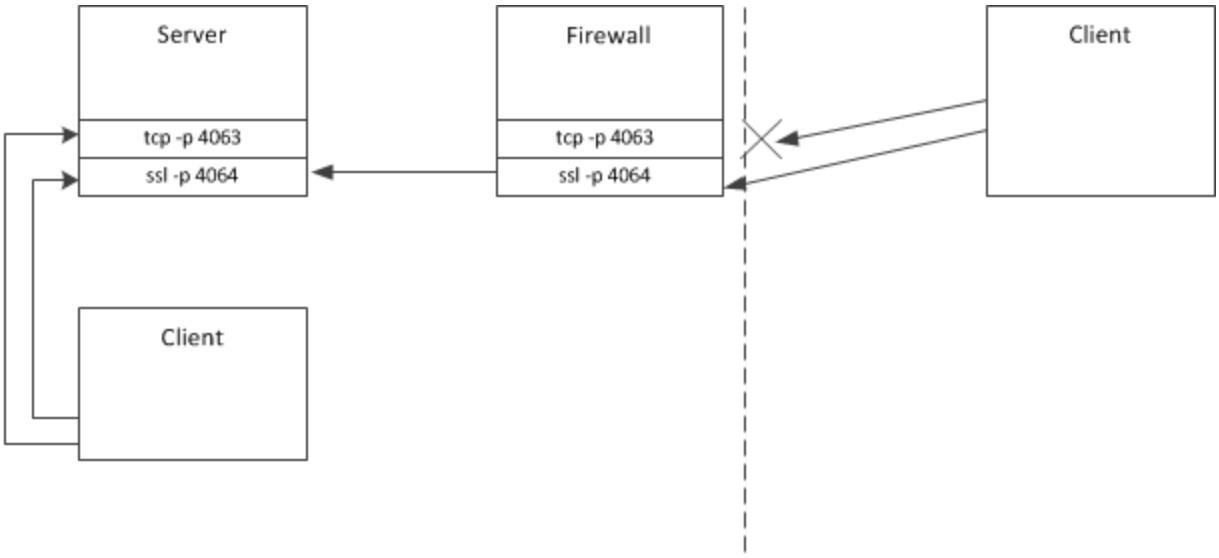

An application of multiple protocol endpoints.

The firewall in the illustration is configured to block external access to TCP port 4063 and to forward connections to port 4064 to the server machine.

One reason for using TCP behind the firewall is that it is more efficient than SSL and requires less administrative work. Of course, this scenario assumes that internal clients can be trusted, which is not true in many environments.

For more information on using SSL in complex network architectures, refer to our discussion of the Glacier2 router.

See Also

- Plug-in Facility
- Configuring IceSSL
- Object Adapter Endpoints
- Proxy Endpoints
- Glacier2
- Ice.Plugin.*
- Ice.InitPlugins
- Ice.PluginLoadOrder

Copyright 2016, ZeroC, Inc.

Configuring IceSSL

After installing IceSSL, an application typically needs to define a handful of additional properties to configure settings such as the location of certificate and key files. This page provides an introduction to configuring the plug-in for each of the supported language mappings.

On this page:

- Configuring IceSSL for C++ and .NET on Windows
  - Using Certificate Files
  - Using Certificate Stores
- Configuring IceSSL for C++ on OS X
  - Keychain Examples for OS X
  - ADH Example for OS X
- Configuring IceSSL for C++ on Linux
  - DSA Example for C++
  - RSA and DSA Example for C++
  - ADH Example for C++
- Configuring IceSSL for Java
  - DSA Example for Java
  - ADH Example for Java
- Configuring IceSSL for Ice Touch
  - Keychains
- Configuring Ciphersuites
  - Configuring Ciphersuites in C++
  - Configuring Ciphersuites in Java
- Configuring Trust Relationships
  - Trusted Peers
  - Verification Depth
- Configuring Secure Proxies
- IceSSL Diagnostics
  - System Logging in Java
  - System Logging in .NET
  - System Logging on OS X

Configuring IceSSL for C++ and .NET on Windows

IceSSL for .NET and C++ supports loading certificates, keys, and trusted CA certificates from regular files. You can also make use of Windows certificate stores.

The ability to load trusted CA certificates from a file in .NET is a new feature added in Ice 3.6.

Using Certificate Files

Our first example demonstrates how to configure IceSSL with certificate files:

```
# For .NET:
Ice.Plugin.IceSSL=IceSSL.dll:IceSSL.PluginFactory
# For C++:
Ice.Plugin.IceSSL=IceSSL:createIceSSL

IceSSL.DefaultDir=/opt/certs
IceSSL.CAs=cacert.pem
IceSSL.CertFile=cert.pfx
IceSSL.Password=password
```

Notice that the value for `Ice.Plugin.IceSSL` differs depending on the language mapping you're using. The `Ice.Plugin` properties also support an alternate syntax that allows you to define both variations in the same configuration file if desired.

The `IceSSL.DefaultDir` property is a convenient way to specify the default location of your certificate files. The two subsequent properties identify the files containing the trusted CA certificate and the certificate with private key, respectively. When these properties specify a relative path name, as shown here, IceSSL attempts to find these files in the directory defined by `IceSSL.DefaultDir`. The file

Copyright 2016, ZeroC, Inc.

specified in `IceSSL.CAs` should normally be in the Privacy Enhanced Mail (PEM) format, whereas the file in `IceSSL.CertFile` must use the Personal Information Exchange (PFX, also known as PKCS#12) format and contain both a certificate and its corresponding private key. The `IceSSL.Password` property specifies the password used to secure the file.

It is a security risk to define a password in a plain text file, such as an Ice configuration file, because anyone who can gain read access to your configuration file can obtain your password. IceSSL also supports alternate ways to supply a password.

Using Certificate Stores

Windows uses certificate stores as the persistent repositories of certificates and keys. Furthermore, there are two distinct sets of certificate stores, one for the current user and another for the local machine.

Managing Certificates with the Microsoft Management Console

On Windows, you can use the Microsoft Management Console (MMC) to browse the contents of the various certificate stores. To start the console, run `MMC.EXE` from a command window, or choose Run from the Start menu and enter `MMC.EXE`.

Once the console is running, you need to install the Certificates "snap-in" by choosing Add/Remove Snap-in from the File menu. Click the Add button, choose Certificates in the popup window and click Add. If you wish to manage certificates for the current user, select My Current Account and click Finish. To manage certificates for the local computer, select Computer Account and click Next, then select Local Computer and click Finish.

When you have finished adding snap-ins, close the Add Standalone Snap-in window and click OK on the Add/Remove Snap-in window. Your Console Root window now contains a tree structure that you can expand to view the available certificate stores. If you have a certificate in a file that you want to add to a store, click on the desired store, then open the Action menu and select All Tasks/Import.

Configuring IceSSL using Certificate Stores

If the program's certificate and private key are already installed in a certificate store, you can select it using the `IceSSL.FindCert` configur ation property as shown in the following example:

```
Ice.Plugin.IceSSL=...
IceSSL.FindCert=subject:"Quote Server"
IceSSL.CertStore=My
IceSSL.CertStoreLocation=LocalMachine
```

An `IceSSL.FindCert` property executes a query in a particular certificate store and selects all of the certificates that match the given criteria. In the example above, the location of the certificate store is `LocalMachine` (as defined by `IceSSL.CertStoreLocation`) and the store's name is `My` (as defined by `IceSSL.CertStore`). When using MMC to browse the certificate stores, this is equivalent to the store "Personal" in the location "Certificates (Local Computer)."

The other legal value for `IceSSL.CertStoreLocation` is `CurrentUser`. The following table shows the valid values for `IceSSL.CertSt ore` and their equivalents in MMC:

Property Value	MMC Name
`AddressBook`	Other People
`AuthRoot`	Third-Party Root Certification Authorities
`CertificateAuthority`	Intermediate Certification Authorities
`Disallowed`	Untrusted Certificates
`My`	Personal
`Root`	Trusted Root Certification Authorities
`TrustedPeople`	Trusted People
`TrustedPublisher`	Trusted Publishers

The search criteria consists of *name:value* pairs that perform case-insensitive comparisons against the fields of each certificate in the

Copyright 2016, ZeroC, Inc.

specified store, and the special property value * selects every certificate in the store. Typically, however, the criteria should select a single certificate. In a server, IceSSL must supply Windows with the certificate that represents the server's identity; if a configuration matches several certificates, IceSSL chooses one (in an undefined manner) and logs a warning to notify you of the situation.

Selecting a certificate from a store is more secure than using a certificate file via the `IceSSL.CertFile` property because it is not necessary to specify a plain-text password. MMC prompts you for the password when initially importing a certificate into a store, so the password is not required when an application uses that certificate to identify itself.

Trusted CA Certificates

We made no mention of trusted CA certificates in the configuration above. For a program to be able to successfully authenticate a certificate it receives from a peer, all trusted CA certificates must either be loaded via the `IceSSL.CAs` property or already be present in the appropriate certificate stores.

When installing a trusted CA certificate, authentication succeeds only when the certificate is installed into one of the following stores:

- Local Computer / Trusted Root Certification Authorities
- Local Computer / Third-Party Root Certification Authorities
- Current User / Trusted Root Certification Authorities

Note that administrative privileges are required when installing a certificate into a Local Computer store.

Configuring IceSSL for C++ on OS X

Our first example shows the properties that are sufficient in many situations:

```
Ice.Plugin.IceSSL=IceSSL:createIceSSL
IceSSL.DefaultDir=/opt/certs
IceSSL.CertFile=cert.pfx
IceSSL.CAs=ca.pem
IceSSL.Password=password
```

The `IceSSL.DefaultDir` property is a convenient way to specify the default location of your certificate files. The properties that follow it define the files containing the program's certificate with private key, and trusted CA certificate, respectively. Finally, the `IceSSL.Password` property specifies the password necessary to open `cert.pfx`.

It is a security risk to define a password in a plain text file, such as an Ice configuration file, because anyone who can gain read access to your configuration file can obtain your password. IceSSL also supports alternate ways to supply a password.

Keychain Examples for OS X

IceSSL imports the certificate specified by `IceSSL.CertFile` into a keychain. IceSSL uses the user's default keychain unless you choose a different one by defining the `IceSSL.Keychain` property:

```
IceSSL.Keychain=Test Keychain
```

If the specified keychain file does not exist, IceSSL will create it. The file name is opened relative to the user's current working directory unless an absolute path name is provided.

The user's default keychain is normally unlocked after logging into the system, so IceSSL doesn't usually require a password to import a certificate into this keychain. However, if your keychain is not unlocked automatically, or if you've selected a different keychain, you can supply a password using the `IceSSL.KeychainPassword` property:

```
IceSSL.KeychainPassword=password
```


It is a security risk to define a password in a plain text file, such as an Ice configuration file, because anyone who can gain read

Copyright 2016, ZeroC, Inc.

access to your configuration file can obtain your password.

If you don't define `IceSSL.KeychainPassword` and a password is required to open the keychain, OS X will prompt the user for the password.

To use a certificate that's already in a keychain, you can omit the `IceSSL.CertFile` and `IceSSL.Password` properties and use `IceSSL.FindCert` instead:

```
Ice.Plugin.IceSSL=IceSSL:createIceSSL
IceSSL.DefaultDir=/opt/certs
IceSSL.CAs=ca.pem
IceSSL.FindCert=Label:Client
```

Here we've instructed IceSSL to locate the certificate in the default keychain labeled `Client` and use that as the program's identity.

ADH Example for OS X

The following example uses ADH (the Anonymous Diffie-Hellman cipher). ADH is not a good choice in most cases because, as its name implies, there is no authentication of the communicating parties, and it is vulnerable to man-in-the-middle attacks. However, it still provides encryption of the session traffic and requires very little administration and therefore may be useful in certain situations. The configuration properties shown below demonstrate how to use ADH:

```
Ice.Plugin.IceSSL=IceSSL:createIceSSL
IceSSL.Ciphers=(DH_anon*)
IceSSL.DHParams=dhparams.der
IceSSL.VerifyPeer=0
```

The `IceSSL.Ciphers` property enables support for ADH. We also recommend setting `IceSSL.DHParams` with the name of a DER-encoded file containing pre-generated DH parameters.

The `IceSSL.VerifyPeer` property changes the plug-in's default behavior with respect to certificate verification. Without this setting, IceSSL rejects a connection if the peer does not supply a certificate (as is the case with ADH).

Configuring IceSSL for C++ on Linux

Our first example shows the properties that are sufficient in many situations:

```
Ice.Plugin.IceSSL=IceSSL:createIceSSL
IceSSL.DefaultDir=/opt/certs
IceSSL.CertFile=cert.pfx
IceSSL.CAs=ca.pem
IceSSL.Password=password
```

The `IceSSL.DefaultDir` property is a convenient way to specify the default location of your certificate and key files. The two properties that follow it define the files containing the program's certificate with private key and trusted CA certificate, respectively. This example assumes the files contain RSA keys, and IceSSL requires the files to use the Privacy Enhanced Mail (PEM) encoding. Finally, the `IceSSL.Password` property specifies the password of the private key.

It is a security risk to define a password in a plain text file, such as an Ice configuration file, because anyone who can gain read access to your configuration file can obtain your password. IceSSL also supports alternate ways to supply a password.

DSA Example for C++

If you used DSA to generate your keys, one additional property is necessary:

Copyright 2016, ZeroC, Inc.

```
Ice.Plugin.IceSSL=IceSSL:createIceSSL
IceSSL.DefaultDir=/opt/certs
IceSSL.CertFile=cert_dsa.pfx
IceSSL.CAs=ca.pem
IceSSL.Password=password
IceSSL.Ciphers=DEFAULT:DSS
```

The `IceSSL.Ciphers` property adds support for DSS authentication to the plug-in's default set of ciphersuites.

RSA and DSA Example for C++

It is also possible to specify certificates and keys for both RSA and DSA by including two filenames in the `IceSSL.CertFile` property. The file names must be separated using the platform's path separator. The example below demonstrates the configuration:

```
Ice.Plugin.IceSSL=IceSSL:createIceSSL
IceSSL.DefaultDir=/opt/certs
IceSSL.CertFile=cert_rsa.pfx:cert_dsa.pfx
IceSSL.CAs=ca.pem
IceSSL.Password=password
IceSSL.Ciphers=DEFAULT:DSS
```

ADH Example for C++

The following example uses ADH (the Anonymous Diffie-Hellman cipher). ADH is not a good choice in most cases because, as its name implies, there is no authentication of the communicating parties, and it is vulnerable to man-in-the-middle attacks. However, it still provides encryption of the session traffic and requires very little administration and therefore may be useful in certain situations. The configuration properties shown below demonstrate how to use ADH:

```
Ice.Plugin.IceSSL=IceSSL:createIceSSL
IceSSL.Ciphers=ADH:!LOW:!MD5:!EXP:!3DES:@STRENGTH
IceSSL.VerifyPeer=0
```

The `IceSSL.Ciphers` property enables support for ADH (which is disabled by default) and eliminates low-strength ciphersuites.

The `IceSSL.VerifyPeer` property changes the plug-in's default behavior with respect to certificate verification. Without this setting, IceSSL rejects a connection if the peer does not supply a certificate (as is the case with ADH).

Configuring IceSSL for Java

IceSSL uses Java's native format for storing keys and certificates: the keystore.

A keystore is represented as a file containing key pairs and associated certificates, and is usually administered using the `keytool` utility supplied with the Java run time. Keystores serve two roles in Java's SSL architecture:

1. A keystore containing a key pair identifies the peer and is usually closely guarded.
2. A keystore containing public certificates represents the identities of trusted peers and can be freely shared. These keystores are also referred to as "truststores" when they are used to store only trusted certificate chains.

A single keystore file can fulfill both of these purposes.

Java supports a pluggable architecture for keystore implementations in which a system property selects a particular implementation as the default keystore type. IceSSL uses the default keystore type unless otherwise specified.

A password is assigned to each key pair in a keystore, as well as to the keystore itself. IceSSL must be provided with the password for the key pair, but the keystore password is optional. If a keystore password is specified, it is used only to verify the keystore's integrity. IceSSL requires that all of the key pairs in a keystore have the same password.

Copyright 2016, ZeroC, Inc.

Our first example shows the properties that are sufficient in many situations:

```
Ice.Plugin.IceSSL=IceSSL.PluginFactory
IceSSL.DefaultDir=/opt/certs
IceSSL.Keystore=keys.jks
IceSSL.Truststore=ca.jks
IceSSL.Password=password
```

IceSSL resolves the filenames defined in its configuration properties as follows:

1. Attempt to open the file as a class loader resource. This is especially useful for deploying applications with special security restrictions, such as applets.
2. Attempt to open the file in the local file system.
3. If `IceSSL.DefaultDir` is defined, prepend its value and try steps 1 and 2 again. The `IceSSL.DefaultDir` property is a convenient way to specify the default location of your keystore and truststore files.

The `IceSSL.Password` property specifies the password of the key pair.

It is a security risk to define a password in a plain text file, such as an Ice configuration file, because anyone who can gain read access to your configuration file can obtain your password. IceSSL also supports alternate ways to supply a password.

DSA Example for Java

Java supports both RSA and DSA keys. No additional properties are necessary to use DSA:

```
Ice.Plugin.IceSSL=IceSSL.PluginFactory
IceSSL.DefaultDir=/opt/certs
IceSSL.Keystore=dsakeys.jks
IceSSL.Truststore=ca.jks
IceSSL.Password=password
```

ADH Example for Java

The following example uses ADH (the Anonymous Diffie-Hellman cipher). ADH is not a good choice in most cases because, as its name implies, there is no authentication of the communicating parties, and it is vulnerable to man-in-the-middle attacks. However, it still provides encryption of the session traffic and requires very little administration and therefore may be useful in certain situations. The configuration properties shown below demonstrate how to use ADH:

```
Ice.Plugin.IceSSL=IceSSL.PluginFactory
IceSSL.DefaultDir=/opt/certs
IceSSL.Ciphers=NONE (DH_anon.*AES)
IceSSL.VerifyPeer=0
```

The `IceSSL.Ciphers` property enables support for ADH, which is disabled by default. Furthermore, this setting enables only those ADH ciphersuites that support AES encryption and eliminates other, lower-strength ciphersuites that may not be supported by recent JVMs.

The `IceSSL.VerifyPeer` property changes the plug-in's default behavior with respect to certificate verification. Without this setting, IceSSL rejects a connection if the peer does not supply a certificate (as is the case with ADH).

Configuring IceSSL for Ice Touch

In Ice Touch, certificate files are loaded from the application's resource bundle. If the application's target platform is OS X, certificate files can also be loaded directly from the file system. Consider the following properties:

Copyright 2016, ZeroC, Inc.

```
IceSSL.DefaultDir=certs
IceSSL.CAs=cacert.der
IceSSL.CertFile=cert.pfx
IceSSL.Password=password
```

The `IceSSL.DefaultDir` property is a convenient way to specify the location of your certificate files. Defining `IceSSL.DefaultDir` means IceSSL searches for certificate files relative to the specified directory. For the properties in the example above, IceSSL composes the pathnames `certs/cacert.der` and `certs/cert.pfx`. If `IceSSL.DefaultDir` is not defined, IceSSL uses the certificate file pathnames exactly as they are supplied.

As mentioned earlier, IceSSL has different semantics for locating certificate files depending on the target platform. For the iPhone and iPhone simulator, IceSSL attempts to open a certificate file in the application's resource bundle as `Resources/`*`DefaultDir/file`* if `IceSSL.DefaultDir` is defined, or as simply `Resources/`*`file`* otherwise. If the target platform is OS X and the certificate file cannot be found in the resource bundle, IceSSL also attempts to open the file in the file system as *`DefaultDir/file`* if a default directory is specified, or as simply *`file`* otherwise.

IceSSL requires that the CA certificate file specified by `IceSSL.CAs` use the DER format. The certificate file in `IceSSL.CertFile` must use the Personal Information Exchange (PFX, also known as PKCS#12) format and contain both a certificate and its corresponding private key. The `IceSSL.Password` property specifies the password used to secure the certificate file.

Keychains

IceSSL imports the certificate specified by `IceSSL.CertFile` into a keychain. IceSSL uses the `login` keychain by default unless you choose a different one by defining the `IceSSL.Keychain` property:

```
IceSSL.Keychain=Test Keychain
```

The `login` keychain is the user's default keychain, which is normally unlocked after logging into the system. IceSSL does not usually require a password to import a certificate into the `login` keychain. However, if your `login` keychain is not unlocked automatically, or if you have selected a different keychain, you can supply a password using the `IceSSL.KeychainPassword` property:

```
IceSSL.KeychainPassword=password
```

Configuring Ciphersuites

A ciphersuite represents a particular combination of encryption, authentication and hashing algorithms. The IceSSL plug-ins for C++ and Java allow you to configure the ciphersuites that their underlying SSL engines are allowed to negotiate during handshaking with a peer. By default, IceSSL uses the underlying engine's default ciphersuites, but you can define a property to customize the list as we demonstrated above with the ADH examples. Normally the default configuration is chosen to eliminate relatively insecure ciphersuites such as ADH, which is the reason it must be explicitly enabled.

Configuring Ciphersuites in C++

The value of the `IceSSL.Ciphers` property is given directly to the platform's native SSL implementation, which determines the allowable ciphersuites.

On Linux, the list of ciphersuites supported by OpenSSL depends on how the OpenSSL distribution was compiled. You can obtain a complete list of the supported ciphersuites using the `openssl` command:

```
$ openssl ciphers
```

This command will likely generate a long list.

In general, we recommend setting `IceSSL.Trace.Security=1` when experimenting with ciphersuite configurations. When executing your program, pay special attention to the log output to verify the ciphersuites that IceSSL has enabled.

Copyright 2016, ZeroC, Inc.

Refer to the platform notes in the `IceSSL.Ciphers` property description for the details about your platform.

Configuring Ciphersuites in Java

IceSSL for Java interprets the value of `IceSSL.Ciphers` as a sequence of expressions that filter the selected ciphersuites using name and pattern matching. If the property is not defined, the Java security provider's default ciphersuites are used. The following table defines the valid expressions that may appear in the property value.

Expression	Description
`NONE`	Disables all ciphersuites. If specified, it must appear first.
`ALL`	Enables all supported ciphersuites. If specified, it must appear first. This expression should be used with caution, as it may enable low-security ciphersuites.
`NAME`	Enables the ciphersuite matching the given name.
`!`*`NAME`*	Disables the ciphersuite matching the given name.
`(EXP)`	Enables ciphersuites whose names contain the regular expression *EXP*.
`!(EXP)`	Disables ciphersuites whose names contain the regular expression *EXP*.

To determine the set of enabled ciphersuites, the plug-in begins with a list of ciphersuite names containing the default set as determined by the security provider. The expressions in the property value add and remove ciphersuites from this list and are evaluated in the order of appearance. For example, consider the following property definition:

```
IceSSL.Ciphers=NONE (RSA.*AES) !(EXPORT)
```

The expressions in this property have the following effects:

- `NONE` clears the list of enabled ciphersuites.
- `(RSA.*AES)` is a regular expression that enables ciphersuites whose names contain the string "RSA" followed by "AES", meaning ciphersuites using RSA authentication and AES encryption.
- `!(EXPORT)` is a regular expression that disables any of the selected ciphersuites whose names contain the string "EXPORT", meaning ciphersuites having export-quality strength.

As another example, this property adds anonymous Diffie-Hellman to the default set of ciphersuites and disables export ciphersuites:

```
IceSSL.Client.Ciphers=(DH_anon) !(EXPORT)
```

Finally, this example selects only one ciphersuite:

```
IceSSL.Client.Ciphers=NONE SSL_RSA_WITH_RC4_128_SHA
```

Configuring Trust Relationships

Declaring that you trust a certificate authority implies that you trust any peer whose certificate was signed directly or indirectly by that certificate authority. It is necessary to use this broad definition of trust in some applications, such as a public Web server. In more controlled environments, it is a good idea to restrict access as much as possible, and IceSSL provides a number of ways for you to do that.

Trusted Peers

After the low-level SSL engine has completed its authentication process, IceSSL can be configured to take additional steps to verify whether a peer should be trusted. The `IceSSL.TrustOnly` family of properties defines a collection of acceptance and rejection filters that IceSSL applies to the distinguished name of a peer's certificate in order to determine whether to allow the connection to proceed. IceSSL permits the connection if the peer's distinguished name matches any of the acceptance filters and does not match any of the rejection filters.

A distinguished name uniquely identifies a person or entity and is generally represented in the following textual form:

Copyright 2016, ZeroC, Inc.

```
C=US, ST=Florida, L=Palm Beach Gardens, O="ZeroC, Inc.", OU=Servers, CN=Quote
Server
```

Suppose we are configuring a client to communicate with the server whose distinguished name is shown above. If we know that the client is allowed to communicate only with this server, we can enforce this rule using the following property:

```
IceSSL.TrustOnly=O="ZeroC, Inc.", OU=Servers, CN=Quote Server
```

With this property in place, IceSSL allows a connection to proceed only if the distinguished name in the server's certificate matches this filter. The property may contain multiple filters, separated by semicolons, if the client needs to communicate with more than one server. Additional variations of the property are also supported.

If the `IceSSL.TrustOnly` properties do not provide the selectivity you require, the next step is to install a custom certificate verifier.

Verification Depth

In order to authenticate a peer, SSL obtains the peer's certificate chain, which includes the peer's certificate as well as that of the root CA. SSL verifies that each certificate in the chain is valid, but there still remains a subtle security risk. Suppose that we have identified a trusted root CA (via its certificate), and a peer has supplied a valid certificate chain signed by our trusted root CA. It is possible for an attacker to obtain a special signing certificate that is signed by our root CA and therefore trusted implicitly. The attacker can use this certificate to sign fraudulent certificates with the goal of masquerading as a trusted peer, presumably for some nefarious purpose.

We could use the `IceSSL.TrustOnly` properties described above in an attempt to defend against such an attack. However, the attacker could easily manufacture a certificate containing a distinguished name that satisfies the trust properties.

If you know that all trusted peers present certificate chains of a certain length, set the property `IceSSL.VerifyDepthMax` so that IceSSL automatically rejects longer chains. The default value of this property is three, therefore you may need to set it to a larger value if you expect peers to present longer chains.

In situations where you cannot make assumptions about the length of a peer's certificate chain, yet you still want to examine the chain before allowing the connection, you should install a custom certificate verifier.

Configuring Secure Proxies

Proxies may contain any combination of secure and insecure endpoints. An application that requires secure communication can guarantee that proxies it manufactures itself, such as those created by calling `stringToProxy`, contain only secure endpoints. However, the application cannot make the same assumption about proxies received as the result of a remote invocation.

The simplest way to guarantee that all proxies use only secure endpoints is to define the `Ice.Override.Secure` configuration property:

```
Ice.Override.Secure=1
```

Setting this property is equivalent to invoking the proxy method `ice_secure(true)` on every proxy. When enabled, attempting to establish a connection using a proxy that does not contain a secure endpoint results in `NoEndpointException`.

If you want the default behavior of proxies to give precedence to secure endpoints, you can set this property instead:

```
Ice.Default.PreferSecure=1
```

Note that proxies may still attempt to establish connections to insecure endpoints, but they try all secure endpoints first. This is equivalent to invoking `ice_preferSecure(true)` on a proxy.

IceSSL Diagnostics

You can use two configuration properties to obtain more information about the plug-in's activities. Setting `IceSSL.Trace.Security=1` ena bles the plug-in's diagnostic output, which includes a variety of messages regarding events such as ciphersuite selection, peer verification

Copyright 2016, ZeroC, Inc.

and trust evaluation. The other property, `Ice.Trace.Network`, determines how much information is logged about network events such as connections and packets. Note that the output generated by `Ice.Trace.Network` also includes other transports such as TCP and UDP.

System Logging in Java

In Java, you can use a system property that displays a great deal of information about SSL certificates and connections, including the ciphersuite that is selected for use by each connection. For example, the following command sets the system property that activates the diagnostics:

```
$ java -Djavax.net.debug=ssl MyProgram
```

System Logging in .NET

Enabling additional tracing output in .NET requires the creation of an XML file such as the one shown below:

XML

```
<?xml version="1.0" encoding="UTF-8" ?>
<configuration>
    <system.diagnostics>
        <trace autoflush="true"/>
        <sources>
            <source name="System.Net">
                <listeners>
                    <add name="System.Net"/>
                </listeners>
            </source>
            <source name="System.Net.Sockets">
                <listeners>
                    <add name="System.Net"/>
                </listeners>
            </source>
            <source name="System.Net.Cache">
                <listeners>
                    <add name="System.Net"/>
                </listeners>
            </source>
        </sources>
        <sharedListeners>
            <add
            name="System.Net"
            type="System.Diagnostics.TextWriterTraceListener"
            initializeData="trace.txt"
            />
        </sharedListeners>
        <switches>
            <add name="System.Net" value="Verbose"/>
            <add name="System.Net.Sockets" value="Verbose"/>
            <add name="System.Net.Cache" value="Verbose"/>
        </switches>
    </system.diagnostics>
</configuration>
```

In this example, the output is stored in the file `trace.txt`. To activate tracing, give the XML file the same name as your executable with a `.`

Copyright 2016, ZeroC, Inc.

`config` extension (such as `server.exe.config`), and place it in the same directory as the executable.

System Logging on OS X

OS X 10.10 (Yosemite) includes a TLS logging facility. To enable it, run the following command in Terminal:

```
$ sudo defaults write /Library/Preferences/com.apple.security SSLDebugScope -bool
true
```

To see the log entries you must enable logging for warnings in syslog:

```
$ sudo syslog -c 0 -w
```

Then run syslog in 'wait' mode:

```
$ syslog -w
```

See Also

- Proxy Methods
- Filtering Proxy Endpoints
- Public Key Infrastructure
- Using IceSSL
- Programming IceSSL
- Advanced IceSSL Topics
- IceSSL.*

Copyright 2016, ZeroC, Inc.

Programming IceSSL

The IceSSL configuration properties are flexible enough to satisfy the requirements of many applications, and IceSSL supports a public API that offers even more functionality for those applications that need it.

Topics

- Programming IceSSL in CPP
- Programming IceSSL in Java
- Programming IceSSL in .NET

See Also

- Configuring IceSSL

Copyright 2016, ZeroC, Inc.

Programming IceSSL in CPP

This page describes the C++ API for the IceSSL plug-in.

On this page:

- The IceSSL Plugin Interface in C++
- Obtaining SSL Connection Information in C++
- Installing a Certificate Verifier in C++
- Using Certificates in C++
- Using Distinguished Names in C++

The IceSSL Plugin Interface in C++

Applications can interact directly with the IceSSL plug-in using the native C++ class `IceSSL::Plugin`. A reference to a `Plugin` object must be obtained from the communicator in which the plug-in is installed:

C++

```
Ice::CommunicatorPtr communicator = // ...
Ice::PluginManagerPtr pluginMgr = communicator->getPluginManager();
Ice::PluginPtr plugin = pluginMgr->getPlugin("IceSSL");
IceSSL::PluginPtr sslPlugin = IceSSL::PluginPtr::dynamicCast(plugin);
```

The `Plugin` class supports the following methods:

C++

```
namespace IceSSL
{
class Plugin : public Ice::Plugin
{
public:
    virtual void setCertificateVerifier(const CertificateVerifierPtr&)
= 0;

    virtual void setPasswordPrompt(const PasswordPromptPtr&) = 0;

    virtual void setContext(SSL_CTX*) = 0;
    virtual SSL_CTX* getContext() = 0;
};
typedef IceUtil::Handle<Plugin> PluginPtr;
}
```

The `setCertificateVerifier` method installs a custom certificate verifier object that the plug-in invokes for each new connection. The `setPasswordPrompt` method provides an alternate way to supply IceSSL with passwords. We discuss certificate verifiers below and revisit the other methods in our discussion of advanced IceSSL programming.

The `setContext` and `getContext` methods are only available on platforms that use OpenSSL and are rarely used in practice. The `IceSSL.InitOpenSSL` property is typically used in conjunction with `setContext`.

Obtaining SSL Connection Information in C++

You can obtain information about any SSL connection using the `getInfo` operation on a `Connection` object. The most-derived type of the

Copyright 2016, ZeroC, Inc.

object returned by `getInfo` varies depending on the transport being used. IceSSL defines two base types in Slice:

Slice

```
module Ice {
    local class ConnectionInfo {
        bool incoming;
        string adapterName;
    };

    local class IPConnectionInfo extends ConnectionInfo {
        string localAddress;
        int localPort;
        string remoteAddress;
        int remotePort;
    };
};

module IceSSL {
    local class ConnectionInfo extends Ice::IPConnectionInfo {
        string cipher;
        Ice::StringSeq certs;
        bool verified;
    };

    local class WSSConnectioninfo extends ConnectionInfo {
        Ice::HeaderDict headers;
    };
};
```

The `cipher` member is a description of the ciphersuite that SSL negotiated for this connection. The local and remote address information is provided in `localAddress` and `remoteAddress`, respectively. The `verified` member indicates whether IceSSL was able to successfully verify the peer's certificate.

The inherited `incoming` member indicates whether the connection is inbound (a server connection) or outbound (a client connection). Finally, if `incoming` is true, the `adapterName` member supplies the name of the object adapter that hosts the endpoint.

The `Ice::ConnectionInfo` object can be narrowed to either `IceSSL::ConnectionInfo` for an SSL endpoint or `IceSSL::WSSConnectionInfo` for a WSS (secure WebSocket) endpoint. Furthermore, IceSSL defines two native C++ subclasses as follows:

Copyright 2016, ZeroC, Inc.

C++

```
namespace IceSSL {
class NativeConnectionInfo : public ConnectionInfo {
public:
    std::vector<CertificatePtr> nativeCerts;
};
typedef IceUtil::Handle<NativeConnectionInfo> NativeConnectionInfoPtr;

class WSSNativeConnectionInfo : public WSSConnectionInfo {
public:
    std::vector<CertificatePtr> nativeCerts;
};
typedef IceUtil::Handle<WSSNativeConnectionInfo>
WSSNativeConnectionInfoPtr;
}
```

The `nativeCerts` member is a vector of certificates representing the peer's certificate chain. The vector is structured so that the first element is the peer's certificate, followed by its signing certificates in the order they appear in the chain, with the root CA certificate as the last element. The vector is empty if the peer did not present a certificate chain.

Installing a Certificate Verifier in C++

A new connection undergoes a series of verification steps before an application is allowed to use it. The low-level SSL engine executes certif icate validation procedures and, assuming the certificate chain is successfully validated, IceSSL performs additional verification as directed by its configuration properties. Finally, if a certificate verifier is installed, IceSSL invokes it to provide the application with an opportunity to decide whether to allow the connection to proceed.

The `CertificateVerifier` interface has only one method:

C++

```
namespace IceSSL
{
class CertificateVerifier : public IceUtil::Shared
{
public:

    virtual bool verify(const NativeConnectionInfoPtr&) = 0;
};
typedef IceUtil::Handle<CertificateVerifier> CertificateVerifierPtr;
}
```

IceSSL rejects the connection if `verify` returns `false`, and allows it to proceed if the method returns `true`. The `verify` method receives a `NativeConnectionInfo` object that describes the connection's attributes.

The following class is a simple implementation of a certificate verifier:

Copyright 2016, ZeroC, Inc.

C++

```
class Verifier : public IceSSL::CertificateVerifier
{
public:

    bool verify(const IceSSL::NativeConnectionInfo& info)
    {
        if (!info.nativeCerts.empty())
        {
            string dn = info.nativeCerts[0].getIssuerDN();
            transform(dn.begin(), dn.end(), dn.begin(), ::tolower);
            if (dn.find("zeroc") != string::npos)
            {
                return true;
            }
        }
        return false;
    }
}
```

In this example, the verifier rejects the connection unless the string `zeroc` is present in the issuer's distinguished name of the peer's certificate. In a more realistic implementation, the application is likely to perform detailed inspection of the certificate chain.

Installing the verifier is a simple matter of calling `setCertificateVerifier` on the plug-in interface:

C++

```
IceSSL::PluginPtr sslPlugin = // ...
sslPlugin->setCertificateVerifier(new Verifier);
```

You should install the verifier before any SSL connections are established.

You can also install a certificate verifier using a custom plug-in to avoid making changes to the code of an existing application.

The Ice run time calls the `verify` method during the connection-establishment process, therefore delays in the `verify` implementation have a direct impact on the performance of the application. Do not make remote invocations from your implementation of `verify`.

Using Certificates in C++

The `ConnectionInfo` class contains a vector of `Certificate` objects representing the peer's certificate chain. `Certificate` is a reference-counted convenience class that hides the complexity of the platform's native SSL API. Its methods are inspired by the Java class `X509Certificate`:

Copyright 2016, ZeroC, Inc.

C++

```
namespace IceSSL
{
class Certificate : public IceUtil::Shared
{
public:

    Certificate(X509CertificateRef);

    static CertificatePtr load(const string&);
    static CertificatePtr decode(const string&);

    bool operator==(const Certificate&) const;
    bool operator!=(const Certificate&) const;

    PublicKeyPtr getPublicKey() const;

    bool verify(const CertificatePtr&) const;

    string encode() const;

    bool checkValidity() const;
    bool checkValidity(const IceUtil::Time&) const;

    IceUtil::Time getNotAfter() const;
    IceUtil::Time getNotBefore() const;

    string getSerialNumber() const;

    DistinguishedName getIssuerDN() const;
    vector<pair<int, string> > getIssuerAlternativeNames();

    DistinguishedName getSubjectDN() const;
    vector<pair<int, string> > getSubjectAlternativeNames();

    int getVersion() const;

    string toString() const;

    X509CertificateRef getCert() const;
};
typedef IceUtil::Handle<Certificate> CertificatePtr;
}
```

The more commonly-used methods are described below; refer to the documentation in `IceSSL/Plugin.h` for information on the methods that are not covered.

The static method `load` creates a certificate from the contents of a certificate file. If an error occurs, the function raises `IceSSL::CertificateReadException`; the reason member provides a description of the problem.

Copyright 2016, ZeroC, Inc.

Use `decode` to obtain a certificate from a PEM-encoded string representing a certificate. The caller must be prepared to catch `IceSSL::CertificateEncodingException` if `decode` fails; the reason member provides a description of the problem.

The `encode` method creates a PEM-encoded string that represents the certificate. The return value can later be passed to `decode` to recreate the certificate.

The `checkValidity` methods determine whether the certificate is valid. The overloading with no arguments returns true if the certificate is valid at the current time; the other overloading accepts an `IceUtil::Time` object and returns true if the certificate is valid at the given time.

The `getNotAfter` and `getNotBefore` methods return instances of `IceUtil::Time` that define the certificate's valid period.

The methods `getIssuerDN` and `getSubjectDN` supply the distinguished names of the certificate's issuer (i.e., the CA that signed the certificate) and subject (i.e., the person or entity to which the certificate was issued). The methods return instances of the class `IceSSL::DistinguishedName`, another convenience class that is described below.

Finally, the `toString` method returns a human-readable string describing the certificate.

Using Distinguished Names in C++

X.509 certificates use a distinguished name to identify a person or entity. The name is an ordered sequence of relative distinguished names that supply values for fields such as common name, organization, state, and country. Distinguished names are commonly displayed in stringified form according to the rules specified by RFC 2253, as shown in the following example:

```
C=US, ST=Florida, L=Palm Beach Gardens, O="ZeroC, Inc.", OU=Servers,
CN=Quote Server
```

`DistinguishedName` is a convenience class provided by IceSSL to simplify the tasks of parsing, formatting and comparing distinguished names.

C++

```
namespace IceSSL
{
class DistinguishedName
{
public:

    DistinguishedName(const std::string&);
    DistinguishedName(const std::list<std::pair<std::string, std::strin
g> >&);

    bool operator==(const DistinguishedName&) const;
    bool operator!=(const DistinguishedName&) const;
    bool operator<(const DistinguishedName&) const;

    bool match(const DistinguishedName&) const;

    operator std::string() const;
};
}
```

The first overloaded constructor accepts a string argument representing a distinguished name encoded using the rules set forth in RFC 2253. The new `DistinguishedName` instance preserves the order of the relative distinguished names in the string. The caller must be prepared to catch `IceSSL::ParseException` if an error occurs during parsing.

Copyright 2016, ZeroC, Inc.

The second overloaded constructor requires a list of type-value pairs representing the relative distinguished names. The new `DistinguishedName` instance preserves the order of the relative distinguished names in the list.

The overloaded operator functions `operator==`, `operator!=`, and `operator<` perform an exact match of distinguished names in which the order of the relative distinguished names is important. For two distinguished names to be equal, they must have the same relative distinguished names in the same order.

The `match` function performs a partial comparison that does not consider the order of relative distinguished names. If `N1` and `N2` are instances of `DistinguishedName`, `N1.match(N2)` returns true if all of the relative distinguished names in `N2` are present in `N1`.

Finally, the string conversion operator encodes the distinguished name in the format described by RFC 2253.

See Also

- Using Connections
- Public Key Infrastructure
- Configuring IceSSL
- Advanced IceSSL Topics

Copyright 2016, ZeroC, Inc.

Programming IceSSL in Java

This page describes the Java API for the IceSSL plug-in.

On this page:

- The IceSSL Plugin Interface in Java
- Obtaining SSL Connection Information in Java
- Installing a Certificate Verifier in Java
- Converting Certificates in Java

The IceSSL Plugin Interface in Java

Applications can interact directly with the IceSSL plug-in using the native Java interface `IceSSL.Plugin`. A reference to a `Plugin` object must be obtained from the communicator in which the plug-in is installed:

Java

```
Ice.Communicator comm = // ...
Ice.PluginManager pluginMgr = comm.getPluginManager();
Ice.Plugin plugin = pluginMgr.getPlugin("IceSSL");
IceSSL.Plugin sslPlugin = (IceSSL.Plugin)plugin;
```

The `Plugin` interface supports the following methods:

Java

```
package IceSSL;

public interface Plugin extends Ice.Plugin
{
    void setContext(javax.net.ssl.SSLContext context);
    javax.net.ssl.SSLContext getContext();

    void setCertificateVerifier(CertificateVerifier verifier);
    CertificateVerifier getCertificateVerifier();

    void setPasswordCallback(PasswordCallback callback);
    PasswordCallback getPasswordCallback();

    void setKeystoreStream(java.io.InputStream stream);

    void setTruststoreStream(java.io.InputStream stream);

    void addSeedStream(java.io.InputStream stream);
}
```

The methods are summarized below:

- `setContext`
  `getContext`
  These methods are for advanced use cases and rarely used in practice.
- `setCertificateVerifier`

Copyright 2016, ZeroC, Inc.

`getCertificateVerifier`
These methods install and retrieve a custom certificate verifier object that the plug-in invokes for each new connection. `getCertificateVerifier` returns null if a verifier has not been set.

- `setPasswordCallback`
  `getPasswordCallback`
  These methods install and retrieve a password callback object that supplies IceSSL with passwords. `getPasswordCallback` returns null if a callback has not been set. Using `setPasswordCallback` is a more secure alternative to setting passwords in clear-text configuration files.
- `setKeystoreStream`
  Supplies an input stream for a keystore containing the key pair. The `IceSSL.Keystore` property is ignored if this method is called with a non-null value. You may supply the same input stream object to this method and to `setTruststoreStream` if your keystore contains your key pair as well as your trusted CA certificates.
- `setTruststoreStream`
  Supplies an input stream for a truststore containing your trusted CA certificates. The `IceSSL.Truststore` property is ignored if this method is called with a non-null value. You may supply the same input stream object to this method and to `setKeystoreStream` if your keystore contains your key pair as well as your trusted CA certificates.
- `addSeedStream`
  Adds an input stream that supplies seed data for the random number generator. You may call this method multiple times if necessary.

Obtaining SSL Connection Information in Java

You can obtain information about any SSL connection using the `getInfo` operation on a `Connection` object. The most-derived type of the object returned by `getInfo` varies depending on the transport being used. IceSSL defines two base types in Slice:

Slice

```
module Ice {
    local class ConnectionInfo {
        bool incoming;
        string adapterName;
    };

    local class IPConnectionInfo extends ConnectionInfo {
        string localAddress;
        int localPort;
        string remoteAddress;
        int remotePort;
    };
};

module IceSSL {
    local class ConnectionInfo extends Ice::IPConnectionInfo {
        string cipher;
        Ice::StringSeq certs;
        bool verified;
    };

    local class WSSConnectioninfo extends ConnectionInfo {
        Ice::HeaderDict headers;
    };
};
```

The `cipher` member is a description of the ciphersuite that SSL negotiated for this connection. The local and remote address information is

Copyright 2016, ZeroC, Inc.

provided in `localAddress` and `remoteAddress`, respectively. The `verified` member indicates whether IceSSL was able to successfully verify the peer's certificate.

The inherited `incoming` member indicates whether the connection is inbound (a server connection) or outbound (a client connection). Finally, if `incoming` is true, the `adapterName` member supplies the name of the object adapter that hosts the endpoint.

The `Ice::ConnectionInfo` object can be narrowed to either `IceSSL::ConnectionInfo` for an SSL endpoint or `IceSSL::WSSConnectionInfo` for a WSS (secure WebSocket) endpoint. Furthermore, IceSSL defines two native Java subclasses as follows:

Java

```
package IceSSL;

public class NativeConnectionInfo extends ConnectionInfo
{
    public java.security.cert.Certificate[] nativeCerts;
}

public class WSSNativeConnectionInfo extends WSSConnectionInfo
{
    public java.security.cert.Certificate[] nativeCerts;
}
```

The `nativeCerts` member is an array of certificates representing the peer's certificate chain. The array is structured so that the first element is the peer's certificate, followed by its signing certificates in the order they appear in the chain, with the root CA certificate as the last element. The array is empty if the peer did not present a certificate chain.

Installing a Certificate Verifier in Java

A new connection undergoes a series of verification steps before an application is allowed to use it. The low-level SSL engine executes certificate validation procedures and, assuming the certificate chain is successfully validated, IceSSL performs additional verification as directed by its configuration properties. Finally, if a certificate verifier is installed, IceSSL invokes it to provide the application with an opportunity to decide whether to allow the connection to proceed.

The `CertificateVerifier` interface has only one method:

Java

```
package IceSSL;

public interface CertificateVerifier
{
    boolean verify(NativeConnectionInfo info);
}
```

IceSSL rejects the connection if `verify` returns `false`, and allows it to proceed if the method returns `true`. The `verify` method receives a `NativeConnectionInfo` object that describes the connection's attributes.

The following class is a simple implementation of a certificate verifier:

Copyright 2016, ZeroC, Inc.

Java

```
import java.security.cert.X509Certificate;
import javax.security.auth.x500.X500Principal;

class Verifier implements IceSSL.CertificateVerifier
{
    public boolean
    verify(IceSSL.NativeConnectionInfo info)
    {
        if (info.nativeCerts != null)
        {
            X509Certificate cert = (X509Certificate)info.nativeCerts[0];
            X500Principal p = cert.getIssuerX500Principal();
            if (p.getName().toLowerCase().indexOf("zeroc") != -1)
            {
                return true;
            }
        }
        return false;
    }
}
```

In this example, the verifier rejects the connection unless the string `zeroc` is present in the issuer's distinguished name of the peer's certificate. In a more realistic implementation, the application is likely to perform detailed inspection of the certificate chain.

Installing the verifier is a simple matter of calling `setCertificateVerifier` on the plug-in interface:

Java

```
IceSSL.Plugin sslPlugin = // ...
sslPlugin.setCertificateVerifier(new Verifier());
```

You should install the verifier before any SSL connections are established. An alternate way of installing the verifier is to define the `IceSSL.CertVerifier` property with the class name of your verifier implementation. IceSSL instantiates the class using its default constructor.

You can also install a certificate verifier using a custom plug-in to avoid making changes to the code of an existing application.

The Ice run time calls the `verify` method during the connection-establishment process, therefore delays in the `verify` implementation have a direct impact on the performance of the application. Do not make remote invocations from your implementation of `verify`.

Converting Certificates in Java

Java does not provide a simple way to create a certificate object from a PEM-encoded string, therefore IceSSL offers the following convenience method:

Copyright 2016, ZeroC, Inc.

Java

```
package IceSSL;

public final class Util
{
    // ...

    public static java.security.cert.X509Certificate
    createCertificate(String certPEM)
        throws java.security.cert.CertificateException;
}
```

Given a string in the PEM format, `createCertificate` returns the equivalent `X509Certificate` object.

See Also

- Using Connections
- Public Key Infrastructure
- Configuring IceSSL
- Advanced IceSSL Topics

Copyright 2016, ZeroC, Inc.

Programming IceSSL in .NET

This page describes the .NET API for the IceSSL plug-in.

On this page:

- The IceSSL Plugin Interface in C#
- Obtaining SSL Connection Information in C#
- Installing a Certificate Verifier in C#
- Converting Certificates in C#

The IceSSL Plugin Interface in C#

Applications can interact directly with the IceSSL plug-in using the native C# interface `IceSSL.Plugin`. A reference to a `Plugin` object must be obtained from the communicator in which the plug-in is installed:

C#

```
Ice.Communicator comm = // ...
Ice.PluginManager pluginMgr = comm.getPluginManager();
Ice.Plugin plugin = pluginMgr.getPlugin("IceSSL");
IceSSL.Plugin sslPlugin = (IceSSL.Plugin)plugin;
```

The `Plugin` interface supports the following methods:

C#

```
namespace IceSSL
{
    using System.Security.Cryptography.X509Certificates;

    abstract public class Plugin : Ice.Plugin
    {
        abstract public void
        setCertificates(X509Certificate2Collection certs);

        abstract public void
        setCACertificates(X509Certificate2Collection certs);
        abstract public void
        setCertificateVerifier(CertificateVerifier verifier);

        abstract public CertificateVerifier
        getCertificateVerifier();

        abstract public void
        setPasswordCallback(PasswordCallback callback);

        abstract public PasswordCallback
        getPasswordCallback();
    }
}
```

The methods are summarized below:

Copyright 2016, ZeroC, Inc.

- `setCertificates`
  `setCACertificates`
  These methods are for advanced use cases and rarely used in practice.
- `setCertificateVerifier`
  `getCertificateVerifier`
  These methods install and retrieve a custom certificate verifier object that the plug-in invokes for each new connection. `getCertificateVerifier` returns null if a verifier has not been set.
- `setPasswordCallback`
  `getPasswordCallback`
  These methods install and retrieve a password callback object that supplies IceSSL with passwords. `getPasswordCallback` returns null if a callback has not been set. Using `setPasswordCallback` is a more secure alternative to setting passwords in clear-text configuration files.

Obtaining SSL Connection Information in C#

You can obtain information about any SSL connection using the `getInfo` operation on a `Connection` object. The most-derived type of the object returned by `getInfo` varies depending on the transport being used. IceSSL defines two base types in Slice:

Slice

```
module Ice {
    local class ConnectionInfo {
        bool incoming;
        string adapterName;
    };

    local class IPConnectionInfo extends ConnectionInfo {
        string localAddress;
        int localPort;
        string remoteAddress;
        int remotePort;
    };
};

module IceSSL {
    local class ConnectionInfo extends Ice::IPConnectionInfo {
        string cipher;
        Ice::StringSeq certs;
        bool verified;
    };

    local class WSSConnectioninfo extends ConnectionInfo {
        Ice::HeaderDict headers;
    };
};
```

The `cipher` member is a description of the ciphersuite that SSL negotiated for this connection. The local and remote address information is provided in `localAddress` and `remoteAddress`, respectively. The `verified` member indicates whether IceSSL was able to successfully verify the peer's certificate.

The inherited `incoming` member indicates whether the connection is inbound (a server connection) or outbound (a client connection). Finally, if `incoming` is true, the `adapterName` member supplies the name of the object adapter that hosts the endpoint.

The `Ice::ConnectionInfo` object can be narrowed to either `IceSSL::ConnectionInfo` for an SSL endpoint or `IceSSL::WSSConnectionInfo` for a WSS (secure WebSocket) endpoint. Furthermore, IceSSL defines two native C# subclasses as follows:

Copyright 2016, ZeroC, Inc.

C#

```
namespace IceSSL
{
    public sealed class NativeConnectionInfo : ConnectionInfo
    {

public System.Security.Cryptography.X509Certificates.X509Certificate2[] nativeCerts
;
    }

    public sealed class WSSNativeConnectionInfo : WSSConnectionInfo
    {
        public System.Security.Cryptography.X509Certificates.X509Certificate2[]
nativeCerts;
    }
}
```

The `nativeCerts` member is an array of certificates representing the peer's certificate chain. The array is structured so that the first element is the peer's certificate, followed by its signing certificates in the order they appear in the chain, with the root CA certificate as the last element. The array is empty if the peer did not present a certificate chain.

Installing a Certificate Verifier in C#

A new connection undergoes a series of verification steps before an application is allowed to use it. The low-level SSL engine executes certif icate validation procedures and, assuming the certificate chain is successfully validated, IceSSL performs additional verification as directed by its configuration properties. Finally, if a certificate verifier is installed, IceSSL invokes it to provide the application with an opportunity to decide whether to allow the connection to proceed.

The `CertificateVerifier` interface has only one method:

C#

```
namespace IceSSL
{
    public interface CertificateVerifier
    {
        bool verify(NativeConnectionInfo info);
    }
}
```

IceSSL rejects the connection if `verify` returns `false`, and allows it to proceed if the method returns `true`. The `verify` method receives a `NativeConnectionInfo` object that describes the connection's attributes.

The following class is a simple implementation of a certificate verifier:

Copyright 2016, ZeroC, Inc.

C#

```
using System.Security.Cryptography.X509Certificates;

class Verifier : IceSSL.CertificateVerifier
{
    public boolean
    verify(IceSSL.NativeConnectionInfo info)
    {
        if (info.nativeCerts != null)
        {
            X500DistinguishedName dn = info.nativeCerts[0].IssuerName;
            if (dn.Name.ToLower().Contains("zeroc"))
            {
                return true;
            }
        }
        return false;
    }
}
```

In this example, the verifier rejects the connection unless the string `zeroc` is present in the issuer's distinguished name of the peer's certificate. In a more realistic implementation, the application is likely to perform detailed inspection of the certificate chain.

Installing the verifier is a simple matter of calling `setCertificateVerifier` on the plug-in interface:

C#

```
IceSSL.Plugin sslPlugin = // ...
sslPlugin.setCertificateVerifier(new Verifier());
```

You should install the verifier before any SSL connections are established. An alternate way of installing the verifier is to define the `IceSSL.CertVerifier` property with the class name of your verifier implementation. IceSSL instantiates the class using its default constructor.

You can also install a certificate verifier using a custom plug-in to avoid making changes to the code of an existing application.

The Ice run time calls the `verify` method during the connection-establishment process, therefore delays in the `verify` implementation have a direct impact on the performance of the application. Do not make remote invocations from your implementation of `verify`.

Converting Certificates in C#

IceSSL offers the following convenience method to create a certificate object from a PEM-encoded string:

Copyright 2016, ZeroC, Inc.

C#

```
namespace IceSSL
{
    using System.Security.Cryptography.X509Certificates;

    public sealed class Util
    {
        // ...

        public static X509Certificate2
        createCertificate(string certPEM);
    }
}
```

Given a string in the PEM format, `createCertificate` returns the equivalent `X509Certificate2` object.

See Also

- Using Connections
- Public Key Infrastructure
- Configuring IceSSL
- Advanced IceSSL Topics

Copyright 2016, ZeroC, Inc.

Advanced IceSSL Topics

This page discusses some additional capabilities of the IceSSL plug-in.

On this page:

- Managing Certificate Passwords
  - Dynamic Properties
  - Password Callbacks in C++
  - Password Callbacks in Java
  - Password Callbacks in .NET
- Manually Configuring IceSSL
- Using Custom Plug-ins with IceSSL

Managing Certificate Passwords

IceSSL may need to obtain a password if it loads a file that contains secure data, such as an encrypted private key. An application can supply a plain-text password in a configuration property, but doing so is a potential security risk. For example, if you define the property on the application's command-line, it may be possible for other users on the same host to see the password simply by obtaining a list of active processes. If you define the property in a configuration file, the password is only as secure as the file in which it is defined.

In highly secure environments where access to a host is tightly restricted, a password can safely be supplied as a plain-text configuration property, or the need for the password can be eliminated altogether by using unsecured key files.

In situations where password security is a concern, the application generally needs to take additional action.

Dynamic Properties

A common technique is to prompt the user for a password and transfer the user's input to a configuration property that the application defines dynamically, as shown below:

C++

```
string password = // ...
Ice::InitializationData initData;
initData.properties = Ice::createProperties(argc, argv);
initData.properties->setProperty("IceSSL.Password", password);
Ice::CommunicatorPtr comm = Ice::initialize(initData);
```

The password must be present in the property set before the communicator is initialized, since IceSSL needs the password during its initialization, and the communicator initializes plug-ins automatically by default.

Password Callbacks in C++

If a password is required but the application has not configured one, IceSSL prompts the user at the terminal during the plug-in's initialization. This behavior is not suitable for some types of applications, such as a program that runs automatically at system startup as a Unix daemon or Windows service.

A terminal prompt is equally undesirable for graphical applications, which would generally prefer to prompt the user in an application window. The dynamic property technique described in the previous section is usually appropriate in this situation.

If your application must supply a password, and you do not want to use a configuration property or a terminal prompt, your remaining option is to install a `PasswordPrompt` object in the plug-in using the `setPasswordPrompt` method. The `PasswordPrompt` class is defined as follows:

Copyright 2016, ZeroC, Inc.

C++

```
namespace IceSSL
{
class PasswordPrompt : public IceUtil::Shared
{
public:

    virtual std::string getPassword() = 0;
};
typedef IceUtil::Handle<PasswordPrompt> PasswordPromptPtr;
}
```

IceSSL invokes `getPassword` on the object when a password is required. If the object returns an incorrect password, IceSSL tries again, up to the limit defined by the `IceSSL.PasswordRetryMax` property.

Note that you must delay the initialization of the IceSSL plug-in until after the `PasswordPrompt` object is installed. To illustrate this point, consider the following example:

C++

```
Ice::CommunicatorPtr communicator = // ...
Ice::PluginManagerPtr pluginMgr = communicator->getPluginManager();
Ice::PluginPtr plugin = pluginMgr->getPlugin("IceSSL");
IceSSL::PluginPtr sslPlugin = IceSSL::PluginPtr::dynamicCast(plugin);
sslPlugin->setPasswordPrompt(new Prompt); // OOPS!
```

This code is incorrect because the `PasswordPrompt` object is installed too late: the communicator is already initialized, which means IceSSL has already attempted to load the file that required a password.

The correct approach is to define the `Ice.InitPlugins` configuration property:

```
Ice.InitPlugins=0
```

This setting causes the communicator to install, but not initialize, its configured plug-ins. The application becomes responsible for initializing the plug-ins, as shown below:

C++

```
Ice::CommunicatorPtr communicator = // ...
Ice::PluginManagerPtr pluginMgr = communicator->getPluginManager();
Ice::PluginPtr plugin = pluginMgr->getPlugin("IceSSL");
IceSSL::PluginPtr sslPlugin = IceSSL::PluginPtr::dynamicCast(plugin);
sslPlugin->setPasswordPrompt(new Prompt);
pluginMgr->initializePlugins();
```


Copyright 2016, ZeroC, Inc.

We assume the communicator was initialized with `Ice.InitPlugins=0`. After installing the `PasswordPrompt` object, the application invokes `initializePlugins` on the plug-in manager to complete the plug-in initialization process.

Password Callbacks in Java

If you do not want to use configuration properties to define passwords, you can install a `PasswordCallback` object in the plug-in using a configuration property, or using the `setPasswordCallback` method. The `PasswordCallback` interface has the following definition:

Java

```
public interface PasswordCallback
{
    char[] getPassword(String alias);
    char[] getTrustStorePassword();
    char[] getKeystorePassword();
}
```

The methods are described below:

- `getPassword`
  Supplies the password for the key with the given alias. The return value must not be null.
- `getTrustStorePassword`
  Supplies the password for a truststore. The method may return null, in which case the integrity of the truststore is not verified.
- `getKeystorePassword` to obtain the password for a keystore.
  Supplies the password for a keystore. The method may return null, in which case the integrity of the keystore is not verified.

For each of these methods, IceSSL clears the contents of the returned array as soon as possible.

The simplest way to install the callback is by defining the configuration property `IceSSL.PasswordCallback`. The property's value is the name of your callback implementation class. IceSSL instantiates the class using its default constructor.

To install the callback manually, you must delay the initialization of the IceSSL plug-in until after the `PasswordCallback` object is installed. To illustrate this point, consider the following example:

Java

```
Ice.Communicator communicator = // ...
Ice.PluginManager pluginMgr = communicator.getPluginManager();
Ice.Plugin plugin = pluginMgr.getPlugin("IceSSL");
IceSSL.Plugin sslPlugin = (IceSSL.Plugin)plugin;
sslPlugin.setPasswordCallback(new CallbackI()); // OOPS!
```

This code is incorrect because the `PasswordCallback` object is installed too late: the communicator is already initialized, which means IceSSL has already attempted to retrieve the certificate that required a password.

The correct approach is to define the `Ice.InitPlugins` configuration property:

```
Ice.InitPlugins=0
```

This setting causes the communicator to install, but not initialize, its configured plug-ins. The application becomes responsible for initializing the plug-ins, as shown below:

Copyright 2016, ZeroC, Inc.

Java

```
Ice.Communicator communicator = // ...
Ice.PluginManager pluginMgr = communicator.getPluginManager();
Ice.Plugin plugin = pluginMgr.getPlugin("IceSSL");
IceSSL.Plugin sslPlugin = (IceSSL.Plugin)plugin;
sslPlugin.setPasswordCallback(new CallbackI());
pluginMgr.initializePlugins();
```

We assume the communicator was initialized with `Ice.InitPlugins=0`. After installing the `PasswordCallback` object, the application invokes `initializePlugins` on the plug-in manager to complete the plug-in initialization process.

Password Callbacks in .NET

If you do not want to use configuration properties to define passwords, you can install a `PasswordCallback` object in the plug-in using a configuration property, or using the `setPasswordCallback` method. The `PasswordCallback` interface has the following definition:

C#

```
using System.Security;

public interface PasswordCallback
{
    SecureString getPassword(string file);
    SecureString getImportPassword(string file);
}
```

The methods are described below:

- `getPassword`
  Supplies the password for the given file. The method may return null if no password is required.
- `getImportPassword`
  Supplies the password for a file from which certificates are imported into the certificate store. The method may return null if no password is required.

The simplest way to install the callback is by defining the configuration property `IceSSL.PasswordCallback`. The property's value is the name of your callback implementation class. IceSSL instantiates the class using its default constructor.

To install the callback manually, you must delay the initialization of the IceSSL plug-in until after the `PasswordCallback` object is installed. To illustrate this point, consider the following example:

C#

```
Ice.Communicator communicator = // ...
Ice.PluginManager pluginMgr = communicator.getPluginManager();
Ice.Plugin plugin = pluginMgr.getPlugin("IceSSL");
IceSSL.Plugin sslPlugin = (IceSSL.Plugin)plugin;
sslPlugin.setPasswordCallback(new CallbackI()); // OOPS!
```


Copyright 2016, ZeroC, Inc.

This code is incorrect because the `PasswordCallback` object is installed too late: the communicator is already initialized, which means IceSSL has already attempted to retrieve the certificate that required a password.

The correct approach is to define the `Ice.InitPlugins` configuration property:

```
Ice.InitPlugins=0
```

This setting causes the communicator to install, but not initialize, its configured plug-ins. The application becomes responsible for initializing the plug-ins, as shown below:

C#

```
Ice.Communicator communicator = // ...
Ice.PluginManager pluginMgr = communicator.getPluginManager();
Ice.Plugin plugin = pluginMgr.getPlugin("IceSSL");
IceSSL.Plugin sslPlugin = (IceSSL.Plugin)plugin;
sslPlugin.setPasswordCallback(new CallbackI());
pluginMgr.initializePlugins();
```

We assume the communicator was initialized with `Ice.InitPlugins=0`. After installing the `PasswordCallback` object, the application invokes `initializePlugins` on the plug-in manager to complete the plug-in initialization process.

Manually Configuring IceSSL

The `Plugin` interface supports a method in each of the supported language mappings that provides an application with more control over the plug-in's configuration.

In C++ with OpenSSL and in Java, an application can call the `setContext` method to supply a pre-configured "context" object used by the underlying SSL engines. In .NET, the `setCertificates` method accepts a collection of certificates that the plug-in should use. In all cases, using one of these methods causes IceSSL to ignore (at a minimum) the configuration properties related to certificates and keys. The application is responsible for accumulating its certificates and keys, and must also deal with any password requirements.

Describing the use of these plug-in methods in detail is outside the scope of this manual, however it is important to understand their prerequisites. In particular, the application needs to have the communicator load the plug-in but not actually initialize it until after the application has had a chance to interact directly with it. (The previous section showed examples of this technique.) The application must define the `Ice.InitPlugins` configuration property:

```
Ice.InitPlugins=0
```

With this setting, the application becomes responsible for completing the plug-in initialization process by invoking `initializePlugins` on the `PluginManager`. The C# example below demonstrates the proper steps:

Copyright 2016, ZeroC, Inc.

C#

```
Ice.Communicator comm = // ...
Ice.PluginManager pluginMgr = comm.getPluginManager();
Ice.Plugin plugin = pluginMgr.getPlugin("IceSSL");
IceSSL.Plugin sslPlugin = (IceSSL.Plugin)plugin;
X509Certificate2Collection certs = // ...
sslPlugin.setCertificates(certs);
pluginMgr.initializePlugins();
```

Using Custom Plug-ins with IceSSL

The Ice plug-in facility is not restricted to protocol implementations. Ice only requires that a plug-in implement the `Ice::Plugin` interface and support the language-specific mechanism for dynamic loading.

The customization options of the IceSSL plug-in make it possible for you to install an application-specific implementation of a certificate verifier in an existing program. For example, you could install a custom certificate verifier in a Glacier2 router without the need to modify Glacier2's source code or rebuild the executable. You would have to write a C++ plug-in to accomplish this, since Glacier2 is written in C++. In short, your plug-in must interact with the IceSSL plug-in and install a certificate verifier.

For this technique to work, it is important that the plug-ins be loaded in a particular order. Specifically, the IceSSL plug-in must be loaded first, followed by the certificate verifier plug-in. By default, Ice loads plug-ins in an undefined order, but you can use the property `Ice.PluginLoadOrder` to specify a particular order.

As an example, let's write a plug-in that installs our simple C++ certificate verifier. Here is the definition of our plug-in class:

C++

```
class VerifierPlugin : public Ice::Plugin
{
public:
    VerifierPlugin(const Ice::CommunicatorPtr & comm) :
        _comm(comm)
    {
    }

    virtual void initialize()
    {
        Ice::PluginManagerPtr pluginMgr = _comm->getPluginManager();
        Ice::PluginPtr plugin = pluginMgr->getPlugin("IceSSL");
        IceSSL::PluginPtr sslPlugin = IceSSL::PluginPtr::dynamicCast(plugin);
        sslPlugin->setCertificateVerifier(new Verifier);
    }

    virtual void destroy()
    {
    }

private:
    Ice::CommunicatorPtr _comm;
};
```


Copyright 2016, ZeroC, Inc.

The class implements the two operations in the `Plugin` interface, `initialize` and `destroy`. The code in `initialize` installs the certificate verifier object, while nothing needs to be done in `destroy`.

The next step is to write the plug-in's factory function, which the communicator invokes to obtain an instance of the plug-in:

C++

```
extern "C"
{

Ice::Plugin*
createVerifierPlugin(
    const Ice::CommunicatorPtr & communicator,
    const string & name,
    const Ice::StringSeq & args)
{
    return new VerifierPlugin(communicator);
}

}
```

We can give the function any name; in this example, we chose `createVerifierPlugin`.

Finally, to install the plug-in we need to define the following properties:

C++

```
Ice.PluginLoadOrder=IceSSL,Verifier
Ice.Plugin.IceSSL=IceSSL:createIceSSL
Ice.Plugin.Verifier=Verifier:createVerifierPlugin
```

The value of `Ice.PluginLoadOrder` guarantees that IceSSL is loaded first. The plug-in specification `Verifier:createVerifierPlugin` identifies the name of the shared library or DLL and the name of the registration function.

There are a few more details you must attend to, such as ensuring that the factory function is exported properly and building the shared library or DLL that contains the new plug-in. Our discussion of the plug-in facility provides more information on developing a plug-in.

See Also

- The Server-Side main Function in C++
- Configuring IceSSL
- Programming IceSSL in CPP
- Programming IceSSL in Java
- Programming IceSSL in .NET
- Plug-in Facility
- IceSSL.*
- Ice.Plugin.*
- Ice.InitPlugins
- Ice.PluginLoadOrder

Copyright 2016, ZeroC, Inc.

Setting up a Certificate Authority

During development, it is convenient to have a simple way of creating new certificates. OpenSSL includes all of the necessary infrastructure for setting up your own certificate authority (CA), but it requires getting more familiar with OpenSSL than is really necessary. To simplify the process, the zeroc-icecertutils PyPi package provides a Python library and and an `iceca` script that allows you to quickly perform the essential tasks:

- initializing a new root CA
- generating new certificates
- converting certificates to match platform-specific requirements.

You are not obligated to use this package; IceSSL accepts certificates from any source as long as they are provided in the appropriate formats. However, you may find this tool sufficient for your development needs, and possibly even for your deployed application as well.

You will find more information on the `IceCertUtils` package and the `iceca` utility on the zeroc-icecertutils PyPI page. You can also take a look at the `makecerts.py` scripts in the `certs` and `cpp/IceGrid/secure` directories from the Ice demos repository for examples on how to use the `IceCertUtils` Python library.

On this page:

See Also

- Configuring IceSSL

Copyright 2016, ZeroC, Inc.

Ice Services

Topics

- Freeze
- FreezeScript
- Glacier2
- IceGrid
- IcePatch2
- IceStorm

Copyright 2016, ZeroC, Inc.

Freeze

Freeze is a collection of services that simplify the use of persistence in Ice applications, as shown below:

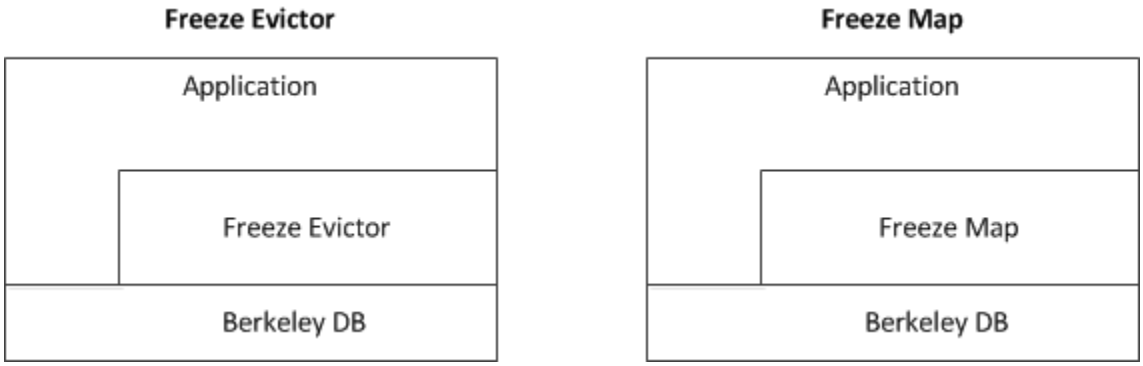

Layer diagram for Freeze persistence services.

The Freeze map is an associative container mapping any Slice key and value types, providing a convenient and familiar interface to a persistent map. Freeze evictors are an especially powerful facility for supporting persistent Ice objects in a highly-scalable implementation.

The Freeze persistence services comprise:

- Freeze evictor
  A highly-scalable implementation of an Ice servant locator that provides automatic persistence and eviction of servants with only minimal application code.
- Freeze map
  A generic associative container. Code generators are provided that produce type-specific maps for Slice key and value types. Applications interact with a Freeze map just like any other associative container, except the keys and values of a Freeze map are persistent.

As you will see from the examples in this discussion, integrating a Freeze map or evictor into your Ice application is quite straightforward: once you define your persistent data in Slice, Freeze manages the mundane details of persistence.

Freeze is implemented using Berkeley DB, a compact and high-performance embedded database. The Freeze map and evictor APIs insulate applications from the Berkeley DB API, but do not prevent applications from interacting directly with Berkeley DB if necessary.

Topics

- Freeze Evictors
- Freeze Maps
- Freeze Catalogs
- Backing Up Freeze Databases

Copyright 2016, ZeroC, Inc.

Freeze Evictors

Freeze evictors combine persistence and scalability features into a single facility that is easily incorporated into Ice applications.

As an implementation of a servant locator, a Freeze evictor takes advantage of the fundamental separation between Ice object and servant to activate servants on demand from persistent storage, and to deactivate them again using customized eviction constraints. Although an application may have thousands of Ice objects in its database, it is not practical to have servants for all of those Ice objects resident in memory simultaneously. The application can conserve resources and gain greater scalability by setting an upper limit on the number of active servants, and letting a Freeze evictor handle the details of servant activation, persistence, and deactivation.

Topics

- Freeze Evictor Concepts
- Background Save Evictor
- Transactional Evictor
- Using a Freeze Evictor in the File System Server

See Also

- Servant Locators

Copyright 2016, ZeroC, Inc.

Freeze Evictor Concepts

This page introduces the Freeze evictor.

On this page:

- Describing Persistent State for an Evictor
- Evictor Servant Semantics
- Evictor Types
- Eviction Strategy
- Detecting Updates to Persistent State
- Iterating an Evictor
- Indexing an Evictor Database
- Using a Servant Initializer
- Application Design Considerations for Evictors

Describing Persistent State for an Evictor

The persistent state of servants managed by a Freeze evictor must be described in Slice. Specifically, every servant must implement a Slice class, and a Freeze evictor automatically stores and retrieves all the (Slice-defined) data members of these Slice classes. Data members that are not specified in Slice are not persistent.

A Freeze evictor relies on the Ice object factory facility to load persistent servants from disk: the evictor creates a brand new servant using the registered factory and then restores the servant's data members. Therefore, for every persistent servant class you define, you need to register a corresponding object factory with the Ice communicator. (For more details on object factories, refer to the C++ mapping or the Java mapping.)

Evictor Servant Semantics

With a Freeze evictor, each <*`object identity, facet`*> pair is associated with its own dedicated persistent object (servant). Such a persistent object cannot serve several identities or facets. Each servant is loaded and saved independently of other servants; in particular, there is no special grouping for the servants that serve the facets of a given Ice object.

Similar to the way you activate servants with an object adapter, the Freeze evictor provides operations named `add`, `addFacet`, `remove`, and `removeFacet`. They have the same signature and semantics, except that with the Freeze evictor, the mapping and the state of the mapped servants is stored in a database.

Evictor Types

Freeze provides two types of evictors with different storage characteristics. The background save evictor records state changes to the database in a background thread, while the transactional evictor records all state changes immediately within the context of a transaction. You can choose the evictor that best fits the persistence requirements of your application.

Eviction Strategy

Both types of evictors associate a queue with their servant map and manage this queue using a "least recently used" eviction algorithm: if the queue is full, the least recently used servant is evicted to make room for a new servant.

Here is the sequence of events for activating a servant as shown in the figure below. Let us assume that we have configured the evictor with a size of five, that the queue is full, and that a request has arrived for a servant that is not currently active. (With a transactional evictor, we also assume this request does not change any persistent state.)

1. A client invokes an operation.
2. The object adapter invokes on the evictor to locate the servant.
3. The evictor first checks its servant map and fails to find the servant, so it instantiates the servant and restores its persistent state from the database.
4. The evictor adds an item for the servant (servant 1) at the head of the queue.
5. The queue's length now exceeds the configured maximum, so the evictor removes servant 6 from the queue as soon as it is eligible for eviction. With a background save evictor, this occurs once there are no outstanding requests pending on servant 6, and once the servant's state has been safely stored in the database. With a transactional save, the servant is removed from the queue immediately.
6. The object adapter dispatches the request to the new servant.

Copyright 2016, ZeroC, Inc.

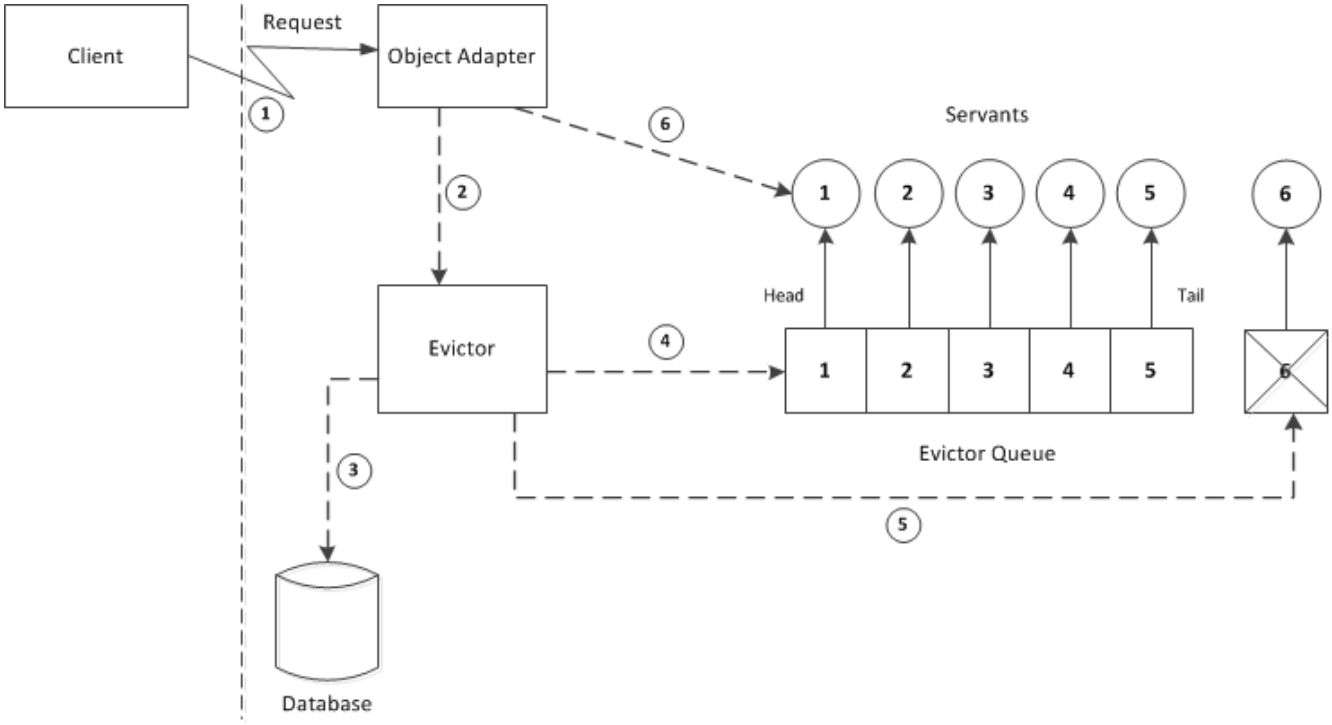

An evictor queue after restoring servant 1 and evicting servant 6.

Detecting Updates to Persistent State

A Freeze evictor considers that a servant's persistent state has been modified when a read-write operation on this servant completes. To indicate whether an operation is read-only or read-write, you add metadata directives to the Slice definitions of the objects:

- The `["freeze:write"]` directive informs the evictor that an operation modifies the persistent state of the target servant.
- The `["freeze:read"]` directive informs the evictor that an operation does not modify the persistent state of the target.

If no metadata directive is present, an operation is assumed to not modify its target.

Here is how you could mark the operations on an interface with these metadata directives:

Slice

```
interface Example {
    ["freeze:read"]  string readonlyOp();
    ["freeze:write"] void   writeOp();
};
```

This marks `readonlyOp` as an operation that does not modify its target, and marks `writeOp` as an operation that does modify its target. Because, without any directive, an operation is assumed to not modify its target, the preceding definition can also be written as follows:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Example {
    string readonlyOp(); // ["freeze:read"] implied
    ["freeze:write"] void writeOp();
};
```

The metadata directives can also be applied to an interface or a class to establish a default. This allows you to mark an interface as `["freeze:write"]` and to only add a `["freeze:read"]` directive to those operations that are read-only, for example:

Slice

```
["freeze:write"]
interface Example {
    ["freeze:read"] string readonlyOp();
                    void   writeOp1();
                    void   writeOp2();
                    void   writeOp3();
};
```

This marks `writeOp1`, `writeOp2`, and `writeOp3` as read-write operations, and `readonlyOp` as a read-only operation.

Note that it is important to correctly mark read-write operations with a `["freeze:write"]` metadata directive — without the directive, Freeze will not know when an object has been modified and may not store the updated persistent state to disk.

Also note that, if you make calls directly on servants (so the calls are not dispatched via the Freeze evictor), the evictor will have no idea when a servant's persistent state is modified; if any such direct call modifies the servant's data members, the update may be lost.

Iterating an Evictor

A Freeze evictor iterator provides the ability to iterate over the identities of the objects stored in an evictor. The operations are similar to Java iterator methods: `hasNext` returns true while there are more elements, and `next` returns the next identity:

Slice

```
local interface EvictorIterator {
    bool hasNext();
    Ice::Identity next();
};
```

You create an iterator by calling `getIterator` on your evictor:

Slice

```
EvictorIterator getIterator(string facet, int batchSize);
```


Copyright 2016, ZeroC, Inc.

The new iterator is specific to a facet (specified by the `facet` parameter). Internally, this iterator will retrieve identities in batches of `batchSize` objects; we recommend using a fairly large batch size to get good performance.

Indexing an Evictor Database

A Freeze evictor supports the use of indexes to quickly find persistent servants using the value of a data member as the search criteria. The types allowed for these indexes are the same as those allowed for Slice dictionary keys.

The `slice2freeze` and `slice2freezej` tools can generate an `Index` class when passed the `--index` option:

`--index` *`CLASS`*`,`*`TYPE`*`,`*`MEMBER`*`[,case-sensitive|case-insensitive]`

`CLASS` is the name of the class to be generated. *`TYPE`* denotes the type of class to be indexed (objects of different classes are not included in this index). *`MEMBER`* is the name of the data member in *`TYPE`* to index. When *`MEMBER`* has type `string`, it is possible to specify whether the index is case-sensitive or not. The default is case-sensitive.

The generated `Index` class supplies three methods whose definitions are mapped from the following Slice operations:

- `sequence<Ice::Identity> findFirst(`*`member-type`* `index, int firstN)`
  Returns up to `firstN` objects of *`TYPE`* whose *`MEMBER`* is equal to `index`. This is useful to avoid running out of memory if the potential number of objects matching the criteria can be very large.
- `sequence<Ice::Identity> find(`*`member-type`* `index)`
  Returns all the objects of *`TYPE`* whose *`MEMBER`* is equal to `index`.
- `int count(`*`member-type`* `index)`
  Returns the number of objects of *`TYPE`* having *`MEMBER`* equal to `index`.

Indexes are associated with a Freeze evictor during evictor creation. See the definition of the `createBackgroundSaveEvictor` and `createTransactionalEvictor` functions for details.

Indexed searches are easy to use and very efficient. However, be aware that an index adds significant write overhead: with Berkeley DB, every update triggers a read from the database to get the old index entry and, if necessary, replace it.

If you add an index to an existing database, by default existing facets are not indexed. If you need to populate a new or empty index using the facets stored in your Freeze evictor, set the property `Freeze.Evictor.`*`env-name.filename`*`.PopulateEmptyIndices` to a non-zero value, which instructs Freeze to iterate over the corresponding facets and create the missing index entries during the call to `createBackgroundSaveEvictor` or `createTransactionalEvictor`. When you use this feature, you must register the object factories for all of the facet types before you create your evictor.

Using a Servant Initializer

In some applications, it may be necessary to initialize a servant after the servant is instantiated by the evictor but before an operation is dispatched to the servant. The Freeze evictor allows an application to specify a servant initializer for this purpose.

To illustrate the sequence of events, let us assume that a request has arrived for a servant that is not currently active:

1. The evictor restores a servant for the target Ice object (and facet) from the database. This involves two steps:
   - The Ice run time locates and invokes the factory for the Ice facet's type, thereby obtaining a new instance with uninitialized data members.
   - The data members are populated from the persistent state.
2. The evictor invokes the application's servant initializer (if any) for the servant.
3. If the evictor is a background-save evictor, it adds the servant to its cache.
4. The evictor dispatches the operation.

With a background-save evictor, the servant initializer is called before the object is inserted into the evictor's internal cache, and *without* holding any internal lock, but in such a way that when the servant initializer is called, the servant is guaranteed to be inserted in the evictor cache.

There is only one restriction on what a servant initializer can do: it must not make a remote invocation on the object (facet) being initialized. Failing to follow this rule will result in deadlocks.

The file system example demonstrates the use of a servant initializer.

Application Design Considerations for Evictors

The Freeze evictor creates a snapshot of a servant's state for persistent storage by marshaling the servant, just as if the servant were being sent "over the wire" as a parameter to a remote invocation. Therefore, the Slice definitions for an object type must include the data members comprising the object's persistent state.

For example, we could define a Slice class as follows:

Copyright 2016, ZeroC, Inc.

Slice

```
class Stateless {
    void calc();
};
```

However, without data members, there will not be any persistent state in the database for objects of this type, and hence there is little value in using the Freeze evictor for this type.

Obviously, Slice object types need to define data members, but there are other design considerations as well. For example, suppose we define a simple application as follows:

Slice

```
class Account {
    ["freeze:write"] void withdraw(int amount);
    ["freeze:write"] void deposit(int amount);

    int balance;
};

interface Bank {
    Account* createAccount();
};
```

In this application, we would use a Freeze evictor to manage `Account` objects that have a data member `balance` representing the persistent state of an account.

From an object-oriented design perspective, there is a glaring problem with these Slice definitions: implementation details (the persistent state) are exposed in the client-server contract. The client cannot directly manipulate the `balance` member because the `Bank` interface returns `Account` proxies, not `Account` instances. However, the presence of the data member may cause unnecessary confusion for client developers.

A better alternative is to clearly separate the persistent state as shown below:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Account {
    ["freeze:write"] void withdraw(int amount);
    ["freeze:write"] void deposit(int amount);
};

interface Bank {
    Account* createAccount();
};

class PersistentAccount implements Account {
    int balance;
};
```

Now the Freeze evictor can manage `PersistentAccount` objects, while clients interact with `Account` proxies. (Ideally, `PersistentAccount` would be defined in a different source file and inside a separate Slice module.)

See Also

- Classes
- C++ Mapping for Classes
- Java Mapping for Classes
- Servant Activation and Deactivation
- Background Save Evictor
- Transactional Evictor
- Dictionaries
- Using a Freeze Evictor in the File System Server

Copyright 2016, ZeroC, Inc.

Background Save Evictor

Freeze provides two types of evictors. This page describes the background save evictor.

Freeze also provides a transactional evictor, with different persistence semantics. The on-disk format of these two types of evictors is the same: you can switch from one type of evictor to the other without any data transformation.

On this page:

- Overview of the Background Save Evictor
- Creating a Background Save Evictor
- The Background Saving Thread
- Synchronization Semantics for the Background Save Evictor
- Preventing Servant Eviction
- Handling Fatal Evictor Errors

Overview of the Background Save Evictor

A background save evictor keeps all its servants in a map and writes the state of newly-created, modified, and deleted servants to disk asynchronously, in a background thread. You can configure how often servants are saved; for example you could decide to save every three minutes, or whenever ten or more servants have been modified. For applications with frequent updates, this allows you to group many updates together to improve performance.

The downside of the background save evictor is recovery from a crash. Because saves are asynchronous, there is no way to force an immediate save to preserve a critical update. Moreover, you cannot group several related updates together: for example, if you transfer funds between two accounts (servants) and a crash occurs shortly after this update, it is possible that, once your application comes back up, you will see the update on one account but not on the other. Your application needs to handle such inconsistencies when restarting after a crash.

Similarly, a background save evictor provides no ordering guarantees for saves. If you update servant 1, servant 2, and then servant 1 again, it is possible that, after recovering from a crash, you will see the latest state for servant 1, but no updates at all for servant 2.

The background save evictor implements the local interface `Freeze::BackgroundSaveEvictor`, which derives from `Freeze::Evictor`.

Creating a Background Save Evictor

You create a background save evictor in C++ with the global function `Freeze::createBackgroundSaveEvictor`, and in Java with the static method `Freeze.Util.createBackgroundSaveEvictor`.

For C++, the signatures are as follows:

Copyright 2016, ZeroC, Inc.

C++

```
BackgroundSaveEvictorPtr
createBackgroundSaveEvictor(
    const ObjectAdapterPtr& adapter,
    const string& envName,
    const string& filename,
    const ServantInitializerPtr& initializer = 0,
    const vector<IndexPtr>& indexes = vector<IndexPtr>(),
    bool createDb = true);

BackgroundSaveEvictorPtr
createBackgroundSaveEvictor(
    const ObjectAdapterPtr& adapter,
    const string& envName,
    DbEnv& dbEnv,
    const string& filename,
    const ServantInitializerPtr& initializer = 0,
    const vector<IndexPtr>& indexes = vector<IndexPtr>(),
    bool createDb = true);
```

For Java, the method signatures are:

Java

```
public static BackgroundSaveEvictor
createBackgroundSaveEvictor(
    Ice.ObjectAdapter adapter,
    String envName,
    String filename,
    ServantInitializer initializer,
    Index[] indexes,
    boolean createDb);

public static BackgroundSaveEvictor
createBackgroundSaveEvictor(
    Ice.ObjectAdapter adapter,
    String envName,
    com.sleepycat.db.Environment dbEnv,
    String filename,
    ServantInitializer initializer,
    Index[] indexes,
    boolean createDb);
```

Both C++ and Java provide two overloaded functions: in one case, Freeze opens and manages the underlying Berkeley DB environment; in the other case, you provide a `DbEnv` object that represents a Berkeley DB environment you opened yourself. (Usually, it is easiest to let Freeze take care of all interactions with Berkeley DB.)

The `envName` parameter represents the name of the underlying Berkeley DB environment, and is also used as the default Berkeley DB home directory. (See `Freeze.DbEnv.`*`env-name`*`.DbHome`.)

Copyright 2016, ZeroC, Inc.

The `filename` parameter represents the name of the Berkeley DB database file associated with this evictor. The persistent state of all your servants is stored in this file.

The `initializer` parameter represents the servant initializer. It is an optional parameter in C++; in Java, pass `null` if you do not need a servant initializer.

The `indexes` parameter is a vector or array of evictor indexes. It is an optional parameter in C++; in Java, pass `null` if your evictor does not define an index.

Finally, the `createDb` parameter tells Freeze what to do when the corresponding Berkeley DB database does not exist. When true, Freeze creates a new database; when false, Freeze raises a `Freeze::DatabaseException`.

The Background Saving Thread

All persistence activity of a background save evictor is handled in a background thread created by the evictor. This thread wakes up periodically and saves the state of all newly-registered, modified, and destroyed servants in the evictor's queue.

For applications that experience bursts of activity that result in a large number of modified servants in a short period of time, you can also configure the evictor's thread to begin saving as soon as the number of modified servants reaches a certain threshold.

Synchronization Semantics for the Background Save Evictor

When the saving thread takes a snapshot of a servant it is about to save, it is necessary to prevent the application from modifying the servant's persistent data members at the same time.

The Freeze evictor and the application need to use a common synchronization to ensure correct behavior. In Java, this common synchronization is the servant itself: the Freeze evictor synchronizes the servant (a Java object) while taking the snapshot. In C++, the servant is required to inherit from the class `IceUtil::AbstractMutex`: the background save evictor locks the servant through this interface while taking a snapshot. On the application side, the servant's implementation is required to use the same mechanism to synchronize all operations that access the servant's Slice-defined data members.

Preventing Servant Eviction

Occasionally, automatically evicting and reloading all servants can be inefficient. You can remove a servant from the evictor's queue by locking this servant "in memory" using the `keep` or `keepFacet` operation on the evictor:

Slice

```
local interface BackgroundSaveEvictor extends Evictor {
    void keep(Ice::Identity id);
    void keepFacet(Ice::Identity id, string facet);
    void release(Ice::Identity id);
    void releaseFacet(Ice::Identity id, string facet);
};
```

`keep` and `keepFacet` are recursive: you need to call `release` or `releaseFacet` for this servant the same number of times to put it back in the evictor queue and make it eligible again for eviction.

Servants kept in memory (using `keep` or `keepFacet`) do not consume a slot in the evictor queue. As a result, the maximum number of servants in memory is approximately the number of kept servants plus the evictor size. (It can be larger while you have many evictable objects that are modified but not yet saved.)

Handling Fatal Evictor Errors

Freeze allows you to register a callback for handling fatal errors encountered by a background save evictor. If no callback is registered, the evictor aborts the application by default.

The application should assume that an evictor will not continue to work properly after encountering a fatal error.

Use `Freeze.Util.registerFatalErrorCallback` (Java) or `Freeze::registerFatalErrorCallback` (C++) to register the

Copyright 2016, ZeroC, Inc.

callback.

For C++, `registerFatalErrorCallback` accepts a function pointer of type `FatalErrorCallback`:

C++

```
typedef void (*FatalErrorCallback)(const BackgroundSaveEvictorPtr&, const
Ice::CommunicatorPtr&);
```

In Java, `registerFatalErrorCallback` accepts a reference to an object that implements `Freeze.FatalErrorCallback`:

Java

```
package Freeze;

public interface FatalErrorCallback
{
    void handleError(Evictor evictor, Ice.Communicator communicator,
RuntimeException ex);
}
```

Note that the `RuntimeException` argument may be `null`.

See Also

- Transactional Evictor
- Freeze Evictor Concepts

Copyright 2016, ZeroC, Inc.

Transactional Evictor

Freeze provides two types of evictors. This page describes the transactional evictor.

Freeze also provides a background save evictor, with different persistence semantics. The on-disk format of these two types of evictors is the same: you can switch from one type of evictor to the other without any data transformation.

On this page:

- Overview of the Transactional Evictor
- Creating a Transactional Evictor
- Read and Write Operations
- Synchronization Semantics for the Transactional Evictor
- Transaction Propagation
- Commit or Rollback on User Exception
- Database Deadlocks and Automatic Retries
- AMD and the Transactional Evictor
- Transactions and Freeze Maps

Overview of the Transactional Evictor

A transactional evictor maintains a servant map, but only keeps read-only servants in this map. The state of these servants corresponds to the latest data on disk. Any servant creation, update, or deletion is performed within a database transaction. This transaction is committed (or rolled back) immediately, typically at the end of the dispatch of the current operation, and the associated servants are then discarded.

With such an evictor, you can ensure that several updates, often on different servants (possibly managed by different transactional evictors) are grouped together: either all or none of these updates occur. In addition, updates are written almost immediately, so crash recovery is a lot simpler: few (if any) updates will be lost, and you can maintain consistency between related persistent objects.

However, an application based on a transactional evictor is likely to write a lot more to disk than an application with a background save evictor, which may have an adverse impact on performance.

Creating a Transactional Evictor

You create a transactional evictor in C++ with the global function `Freeze::createTransactionalEvictor`, and in Java with the static method `Freeze.Util.createTransactionalEvictor`.

For C++, the signatures are as follows:

Copyright 2016, ZeroC, Inc.

C++

```
typedef map<string, string> FacetTypeMap;

TransactionalEvictorPtr
createTransactionalEvictor(
    const ObjectAdapterPtr& adapter,
    const string& envName,
    const string& filename,
    const FacetTypeMap& facetTypes = FacetTypeMap(),
    const ServantInitializerPtr& initializer = 0,
    const vector<IndexPtr>& indexes = vector<IndexPtr>(),
    bool createDb = true);

TransactionalEvictorPtr
createTransactionalEvictor(
    const ObjectAdapterPtr& adapter,
    const string& envName,
    DbEnv& dbEnv,
    const string& filename,
    const FacetTypeMap& facetTypes = FacetTypeMap(),
    const ServantInitializerPtr& initializer = 0,
    const vector<IndexPtr>& indexes = vector<IndexPtr>(),
    bool createDb = true);
```

For Java, the method signatures are:

Java

```
public static TransactionalEvictor
createTransactionalEvictor(
    Ice.ObjectAdapter adapter,
    String envName,
    String filename,
    java.util.Map facetTypes,
    ServantInitializer initializer,
    Index[] indexes,
    boolean createDb);

public static TransactionalEvictor
createTransactionalEvictor(
    Ice.ObjectAdapter adapter,
    String envName,
    com.sleepycat.db.Environment dbEnv,
    String filename,
    java.util.Map facetTypes,
    ServantInitializer initializer,
    Index[] indexes,
    boolean createDb);
```


Copyright 2016, ZeroC, Inc.

Both C++ and Java provide two overloaded functions: in one case, Freeze opens and manages the underlying Berkeley DB environment; in the other case, you provide a `DbEnv` object that represents a Berkeley DB environment you opened yourself. (Usually, it is easier to let Freeze take care of all interactions with Berkeley DB.)

The `envName` parameter represents the name of the underlying Berkeley DB environment, and is also used as the default Berkeley DB home directory. (See `Freeze.DbEnv.`*`env-name`*`.DbHome`.)

The `filename` parameter represents the name of the Berkeley DB database file associated with this evictor. The persistent state of all your servants is stored in this file.

The `facetTypes` parameter allows you to specify a single class type (Slice type ID string) for each facet in your new evictor (see below). Most applications use only the default facet, represented by an empty string. This parameter is optional in C++; in Java, pass `null` if you do not want to specify such a facet-to-type mapping.

The `initializer` parameter represents the servant initializer. It is an optional parameter in C++; in Java, pass `null` if you do not need a servant initializer.

The `indexes` parameter is a vector or array of evictor indexes. It is an optional parameter in C++; in Java, pass `null` if your evictor does not define an index.

Finally, the `createDb` parameter tells Freeze what to do when the corresponding Berkeley DB database does not exist. When true, Freeze creates a new database; when false, Freeze raises a `Freeze::DatabaseException`.

Read and Write Operations

When a transactional evictor processes an incoming request without an associated transaction, it first needs to find out whether the corresponding operation is read-only or read-write (as specified by the `"freeze:read"` and `"freeze:write"` operation metadata). This is straightforward if the evictor knows the target's type; in this case, it simply instantiates and keeps a "dummy" servant to look up the attributes of each operation.

However, if the target type can vary, the evictor needs to look up and sometimes load a read-only servant to find this information. For read-write requests, it will then load the servant from disk a second time (within a new transaction). Once the transaction commits, the read-only servant — sometimes freshly loaded from disk — is discarded.

When you create a transactional evictor with `createTransactionalEvictor`, you can pass a facet name to type ID map to associate a single servant type with each facet and speed up the retrieval of these operation attributes.

Synchronization Semantics for the Transactional Evictor

With a transactional evictor, there is no need to perform any synchronization on the servants managed by the evictor:

- For read-only operations, the application must not modify any data member, and hence there is no need to synchronize. (Many threads can safely read the same data members concurrently.)
- For read-write operations, each operation dispatch gets its own private servant or servants (see transaction propagation below).

Not having to worry about synchronization can dramatically simplify your application code.

Transaction Propagation

Without a distributed transaction service, it is not possible to invoke several remote operations within the same transaction. Nevertheless, Freeze supports transaction propagation for collocated calls: when a request is dispatched within a transaction, the transaction is associated with the dispatch thread and will propagate to any other servant reached through a collocated call. If the target of a collocated call is managed by a transactional evictor associated with the same database environment, Freeze reuses the propagated transaction to load the servant and dispatch the request. This allows you to group updates to several servants within a single transaction.

You can also control how a transactional evictor handles an incoming transaction through optional metadata added after `"freeze:write"` and `"freeze:read"`. There are six valid directives:

- `freeze:read:never`
  Verify that no transaction is propagated to this operation. If a transaction is present, the transactional evictor raises a `Freeze::DatabaseException`.
- `freeze:read:supports`
  Accept requests with or without a transaction, and re-use the transaction if present. `"supports"` is the default for `"freeze:read"` operations.
- `freeze:read:mandatory` and `freeze:write:mandatory`
  Verify that a transaction is propagated to this operation. If there is no transaction, the transactional evictor raises a `Freeze::Datab`

Copyright 2016, ZeroC, Inc.

`aseException`.

- `freeze:read:required` and `freeze:write:required`
  Accept requests with or without a transaction, and re-use the transaction if present. If no transaction is propagated, the transactional evictor creates a new transaction before dispatching the request. `"required"` is the default for `"freeze:write"` operations.

Commit or Rollback on User Exception

When a transactional evictor processes an incoming read-write request, it starts a new database transaction, loads a servant within the transaction, dispatches the request, and then either commits or rolls back the transaction depending on the outcome of this dispatch. If the dispatch does not raise an exception, the transaction is committed just before the response is sent back to the client. If the dispatch raises a system exception, the transaction is rolled back. If the dispatch raises a user exception, by default, the transaction is committed. However, you can configure Freeze to rollback on user-exceptions by setting `Freeze.Evictor.`*`env-name.fileName`*`.RollbackOnUserExcepti on` to a non-zero value.

Database Deadlocks and Automatic Retries

When reading and writing in separate concurrent transactions, deadlocks are likely to occur. For example, one transaction may lock pages in a particular order while another transaction locks the same pages in a different order; the outcome is a deadlock. Berkeley DB automatically detects such deadlocks, and "kills" one of the transactions.

With a Freeze transactional evictor, the application does not need to catch any deadlock exceptions or retry when deadlock occurs because the transactional evictor automatically retries its transactions whenever it encounters a deadlock situation.

However, this can affect how you implement your operations: for any operation called within a transaction (mainly read-write operations), you must anticipate the possibility of several calls for the same request, all in the same dispatch thread.

AMD and the Transactional Evictor

When a transactional evictor dispatches a read-write operation implemented using AMD, it starts a transaction before dispatching the request, and commits or rolls back the transaction when the dispatch is done. Two threads are involved here: the *dispatch thread* and the *cal lback thread*. The dispatch thread is a thread from an Ice thread pool tasked with dispatching a request, and the callback thread is the thread that invokes the AMD callback to send the response to the client. These threads may be one and the same if the servant invokes the AMD callback from the dispatch thread.

It is important to understand the threading semantics of an AMD request with respect to the transaction:

- If a successful AMD response is sent from the dispatch thread, the transaction is committed *after* the response is sent. If a deadlock occurs during this commit, the request is not retried and the client receives no indication of the failure.
- If a successful AMD response is sent from another thread, the evictor commits its transaction when the dispatch thread completes, regardless of whether the servant has sent the AMD response. The callback thread waits until the transaction has been committed by the dispatch thread before sending the response.
- If a commit results in a deadlock and the AMD response has not yet been sent, the evictor cancels the original AMD callback and automatically retries the request again with a new AMD callback. Invocations on the original AMD callback are ignored (`ice_respo nse` and `ice_exception` on this callback do nothing).
- Otherwise, if the servant sends an exception via the AMD callback, the response is sent directly to the client.

Transactions and Freeze Maps

A transactional evictor uses the same transaction objects as Freeze maps, which allows you to update a Freeze map within a transaction managed by a transactional evictor.

You can get the current transaction created by a transactional evictor by calling `getCurrentTransaction`. Then, you would typically retrieve the associated Freeze connection (with `getConnection`) and construct a Freeze map using this connection:

Slice

```
local interface TransactionalEvictor extends Evictor {
    Transaction getCurrentTransaction();
    void setCurrentTransaction(Transaction tx);
};
```


Copyright 2016, ZeroC, Inc.

A transactional evictor also gives you the ability to associate your own transaction with the current thread, using `setCurrentTransaction`. This is useful if you want to perform many updates within a single transaction, for example to add or remove many servants in the evictor. (A less convenient alternative is to implement all such updates within a read-write operation on some object.)

See Also

- Background Save Evictor
- Type IDs
- Freeze Evictor Concepts
- Freeze Maps

Copyright 2016, ZeroC, Inc.

Using a Freeze Evictor in the File System Server

In this section, we present file system implementations that use a transactional evictor. The implementations are based on the ones discussed in Object Life Cycle, and in this section we only discuss code that illustrates use of the Freeze evictor.

In general, incorporating a Freeze evictor into your application requires the following steps:

1. Evaluate your existing Slice definitions for a suitable persistent object type.
2. If no suitable type is found, you typically define a new derived class that captures your persistent state requirements. Consider placing these definitions in a separate file: they are only used by the server for persistence, and therefore do not need to appear in the "public" definitions required by clients. Also consider placing your persistent types in a separate module to avoid name clashes.
3. If you use indexes with your evictor, generate code (using `slice2freeze` or `slice2freezej`) for your new definitions.
4. Create an evictor and register it as a servant locator with an object adapter.
5. Create instances of your persistent type and register them with the evictor.

Persistent Types for File System Evictor

Fortunately, it is unnecessary for us to change any of the existing file system Slice definitions to incorporate the Freeze evictor. However, we do need to add metadata definitions to inform the evictor which operations modify object state:

Slice

```
module Filesystem {
    // ...

    interface Node {
        idempotent string name();

        ["freeze:write"]
        void destroy() throws PermissionDenied;
    };

    // ...

    interface File extends Node {
        idempotent Lines read();

        ["freeze:write"]
        idempotent void write(Lines text) throws GenericError;
    };

    // ...

    interface Directory extends Node {
        idempotent NodeDescSeq list();

        idempotent NodeDesc find(string name) throws NoSuchName;

        ["freeze:write"]
        File* createFile(string name) throws NameInUse;

        ["freeze:write"]
        Directory* createDirectory(string name) throws NameInUse;
    };
};
```


Copyright 2016, ZeroC, Inc.

These definitions are identical to the original ones, with the exception of the added `["freeze:write"]` directives.

The remaining definitions are in derived classes:

Slice

```
#include <Filesystem.ice>

module Filesystem {
    class PersistentDirectory;

    class PersistentNode implements Node {
        string nodeName;
        PersistentDirectory* parent;
    };

    class PersistentFile extends PersistentNode implements File {
        Lines text;
    };

    dictionary<string, NodeDesc> NodeDict;

    class PersistentDirectory extends PersistentNode implements Directory {
        ["freeze:write"]
        void removeNode(string name);

        NodeDict nodes;
    };
};
```

As you can see, we have sub-classed all of the file system interfaces. Let us examine each one in turn.

The `PersistentNode` class adds two data members: `nodeName` and `parent`.

We used `nodeName` instead of `name` because `name` is already used as an operation in the `Node` interface.

The file system implementation requires that a child node knows its parent node in order to properly implement the `destroy` operation. Previous implementations had a state member of type `DirectoryI`, but that is not workable here. It is no longer possible to pass the parent node to the child node's constructor because the evictor may be instantiating the child node (via a factory), and the parent node will not be known. Even if it were known, another factor to consider is that there is no guarantee that the parent node will be active when the child invokes on it, because the evictor may have evicted it. We solve these issues by storing a proxy to the parent node. If the child node invokes on the parent node via the proxy, the evictor automatically activates the parent node if necessary.

The `PersistentFile` class is very straightforward, simply adding a `text` member representing the contents of the file. Notice that the class extends `PersistentNode`, and therefore inherits the state members declared by the base class.

Finally, the `PersistentDirectory` class defines the `removeNode` operation, and adds the `nodes` state member representing the immediate children of the directory node. Since a child node contains only a proxy for its `PersistentDirectory` parent, and not a reference to an implementation class, there must be a Slice-defined operation that can be invoked when the child is destroyed.

If we had followed our earlier advice, we would have defined `Node`, `File`, and `Directory` classes in a separate `PersistentFilesystem` module, but in this example we use the existing `Filesystem` module for the sake of simplicity.

Topics

- Adding an Evictor to the C++ File System Server
- Adding an Evictor to the Java File System Server

See Also

Copyright 2016, ZeroC, Inc.

- Object Life Cycle
- Freeze Evictors
- Freeze Evictor Concepts

Copyright 2016, ZeroC, Inc.

Adding an Evictor to the C++ File System Server

On this page:

- The Server main Program in C++
- The Persistent Servant Class Definitions in C++
- Implementing a Persistent FileI in C++
- Implementing a Persistent DirectoryI in C++
- Implementing NodeFactory in C++

The Server `main` Program in C++

The server's `main` program is responsible for creating the evictor and initializing the root directory node. Many of the administrative duties, such as creating and destroying a communicator, are handled by the `Ice::Application` class. Our server `main` program has now become the following:

C++

```
#include <PersistentFilesystemI.h>

using namespace std;
using namespace Filesystem;

class FilesystemApp : virtual public Ice::Application
{
public:

    FilesystemApp(const string& envName) :
        _envName(envName)
    {
    }

    virtual int run(int, char*[])
    {
        Ice::ObjectFactoryPtr factory = new NodeFactory;
        communicator()->addObjectFactory(factory, PersistentFile::ice_staticId());
        communicator()->addObjectFactory(factory,
PersistentDirectory::ice_staticId());

        Ice::ObjectAdapterPtr adapter =
communicator()->createObjectAdapter("EvictorFilesystem");

        Freeze::EvictorPtr evictor =
            Freeze::createTransactionalEvictor(adapter, _envName, "evictorfs");
        FileI::_evictor = evictor;
        DirectoryI::_evictor = evictor;

        adapter->addServantLocator(evictor, "");

        Ice::Identity rootId;
        rootId.name = "RootDir";
        if(!evictor->hasObject(rootId))
        {
            PersistentDirectoryPtr root = new DirectoryI;
            root->nodeName = "/";
            evictor->add(root, rootId);
        }
```


Copyright 2016, ZeroC, Inc.

```
        adapter->activate();

        communicator()->waitForShutdown();
        if(interrupted())
        {
            cerr << appName() << ": received signal, shutting down" << endl;
        }

        return 0;
    }

private:

    string _envName;
};

int
main(int argc, char* argv[])
{
    FilesystemApp app("db");
```


Copyright 2016, ZeroC, Inc.

```
    return app.main(argc, argv, "config.server");
}
```

Let us examine the changes in detail. First, we are now including `PersistentFilesystemI.h`. This header file includes all of the other Freeze (and Ice) header files this source file requires.

Next, we define the class `FilesystemApp` as a subclass of `Ice::Application`, and provide a constructor taking a string argument:

C++

```
    FilesystemApp(const string& envName) :
        _envName(envName) { }
```

The string argument represents the name of the database environment, and is saved for later use in `run`.

One of the first tasks `run` performs is installing the Ice object factories for `PersistentFile` and `PersistentDirectory`. Although these classes are not exchanged via Slice operations, they are marshalled and unmarshalled in exactly the same way when saved to and loaded from the database, therefore factories are required. A single instance of `NodeFactory` is installed for both types:

C++

```
        Ice::ObjectFactoryPtr factory = new NodeFactory;
        communicator()->addObjectFactory(factory, PersistentFile::ice_staticId());
        communicator()->addObjectFactory(factory,
PersistentDirectory::ice_staticId());
```

After creating the object adapter, the program initializes a transactional evictor by invoking `createTransactionalEvictor`. The third argument to `createTransactionalEvictor` is the name of the database file, which by default is created if it does not exist. The new evictor is then added to the object adapter as a servant locator for the default category:

C++

```
        NodeI::_evictor = Freeze::createTransactionalEvictor(adapter, _envName,
"evictorfs");
        adapter->addServantLocator(NodeI::_evictor, "");
```

Next, the program creates the root directory node if it is not already being managed by the evictor:

Copyright 2016, ZeroC, Inc.

C++

```
        Ice::Identity rootId;
        rootId.name = "RootDir";
        if(!evictor->hasObject(rootId))
        {
            PersistentDirectoryPtr root = new DirectoryI;
            root->nodeName = "/";
            evictor->add(root, rootId);
        }
```

Finally, the `main` function instantiates the `FilesystemApp`, passing `db` as the name of the database environment:

C++

```
int
main(int argc, char* argv[])
{
    FilesystemApp app("db");
    return app.main(argc, argv, "config.server");
}
```

The Persistent Servant Class Definitions in C++

The servant classes must also be changed to incorporate the Freeze evictor. We no longer derive the servants from a common base class. Instead, `FileI` and `DirectoryI` each have their own `_destroyed` and `_mutex` members, as well as a static `_evictor` smart pointer that points at the transactional evictor:

Copyright 2016, ZeroC, Inc.

C++

```
#include <PersistentFilesystem.h>
#include <IceUtil/IceUtil.h>
#include <Freeze/Freeze.h>

namespace Filesystem {

class FileI : virtual public PersistentFile {
public:

    FileI();

    // Slice operations...

    static Freeze::EvictorPtr _evictor;

private:

    bool _destroyed;
    IceUtil::Mutex _mutex;
};

class DirectoryI : virtual public PersistentDirectory {
public:

    DirectoryI();

    // Slice operations...

    virtual void removeNode(const std::string&, const Ice::Current&);

    static Freeze::EvictorPtr _evictor;

public:
    bool _destroyed;
    IceUtil::Mutex _mutex;
};
```

In addition to the node implementation classes, we have also declared an object factory:

C++

```
namespace Filesystem {
    class NodeFactory : virtual public Ice::ObjectFactory {
    public:
        virtual Ice::ObjectPtr create(const std::string&);
        virtual void destroy();
    };
```


Copyright 2016, ZeroC, Inc.

Implementing a Persistent `FileI` in C++

The `FileI` methods are mostly trivial, because the Freeze evictor handles persistence for us:

C++

```
Filesystem::FileI::FileI() : _destroyed(false)
{
}

string
Filesystem::FileI::name(const Ice::Current& c)
{
    IceUtil::Mutex::Lock lock(_mutex);

    if (_destroyed) {
        throw Ice::ObjectNotExistException(__FILE__, __LINE__, c.id, c.facet, c.ope
ration);
    }

    return nodeName;
}

void
Filesystem::FileI::destroy(const Ice::Current& c)
{
    {
        IceUtil::Mutex::Lock lock(_mutex);

        if (_destroyed) {
            throw Ice::ObjectNotExistException(__FILE__, __LINE__, c.id, c.facet, c
.operation);
        }
        _destroyed = true;
    }

    //
    // Because we use a transactional evictor,
    // these updates are guaranteed to be atomic.
    //
    parent?>removeNode(nodeName);
    _evictor?>remove(c.id);
}

Filesystem::Lines
Filesystem::FileI::read(const Ice::Current& c)
{
    IceUtil::Mutex::Lock lock(_mutex);

    if (_destroyed) {
        throw Ice::ObjectNotExistException(__FILE__, __LINE__, c.id, c.facet, c.ope
ration);
    }
```


Copyright 2016, ZeroC, Inc.

```
    return text;
}

void
Filesystem::FileI::write(const Filesystem::Lines& text,
                         const Ice::Current& c)
{
    IceUtil::Mutex::Lock lock(_mutex);

    if (_destroyed) {
        throw Ice::ObjectNotExistException(__FILE__, __LINE__, c.id, c.facet, c.ope
ration);
    }
```


Copyright 2016, ZeroC, Inc.

```
    this?>text = text;
}
```

The code checks that the node has not been destroyed before acting on the invocation by updating or returning state. Note that `destroy` must update two separate nodes: as well as removing itself from the evictor, the node must also update the parent's node map. Because we are using a transactional evictor, the two updates are guaranteed to be atomic, so it is impossible to the leave the file system in an inconsistent state.

Implementing a Persistent `DirectoryI` in C++

The `DirectoryI` implementation requires more substantial changes. We begin our discussion with the `createDirectory` operation:

C++

```
Filesystem::DirectoryPrx
Filesystem::DirectoryI::createDirectory(const string& name, const Ice::Current& c)
{
    IceUtil::Mutex::Lock lock(_mutex);

    if (_destroyed) {
        throw Ice::ObjectNotExistException(__FILE__, __LINE__, c.id, c.facet, c.ope
ration);
    }

    if (name.empty() || nodes.find(name) != nodes.end()) {
        throw NameInUse(name);
    }

    Ice::Identity id;
    id.name = IceUtil::generateUUID();
    PersistentDirectoryPtr dir = new DirectoryI;
    dir?>nodeName = name;
    dir?>parent = PersistentDirectoryPrx::uncheckedCast(c.adapter?>createProxy(c.id
));
    DirectoryPrx proxy = DirectoryPrx::uncheckedCast(_evictor?>add(dir, id));

    NodeDesc nd;
    nd.name = name;
    nd.type = DirType;
    nd.proxy = proxy;
    nodes[name] = nd;

    return proxy;
}
```

After validating the node name, the operation obtains a unique identity for the child directory, instantiates the servant, and registers it with the Freeze evictor. Finally, the operation creates a proxy for the child and adds the child to its node table.

The implementation of the `createFile` operation has the same structure as `createDirectory`:

Copyright 2016, ZeroC, Inc.

C++

```
Filesystem::FilePrx
Filesystem::DirectoryI::createFile(const string& name,
                                   const Ice::Current& c)
{
    IceUtil::Mutex::Lock lock(_mutex);

    if (_destroyed) {
        throw Ice::ObjectNotExistException(__FILE__, __LINE__, c.id, c.facet, c.ope
ration);
    }

    if (name.empty() || nodes.find(name) != nodes.end()) {
        throw NameInUse(name);
    }

    Ice::Identity id;
    id.name = IceUtil::generateUUID();
    PersistentFilePtr file = new FileI;
    file?>nodeName = name;
    file?>parent = PersistentDirectoryPrx::uncheckedCast(c.adapter?>createProxy(c.i
d));
    FilePrx proxy = FilePrx::uncheckedCast(_evictor?>add(file, id));

    NodeDesc nd;
    nd.name = name;
    nd.type = FileType;
    nd.proxy = proxy;
    nodes[name] = nd;

    return proxy;
}
```

Implementing `NodeFactory` in C++

We use a single factory implementation for creating two types of Ice objects: `PersistentFile` and `PersistentDirectory`. These are the only two types that the Freeze evictor will be restoring from its database.

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPtr
Filesystem::NodeFactory::create(const string& type)
{
    if (type == PersistentFile::ice_staticId())
        return new FileI;
    else if (type == PersistentDirectory::ice_staticId())
        return new DirectoryI;
    else {
        assert(false);
        return 0;
    }
}

void
Filesystem::NodeFactory::destroy()
{
}
```

The remaining Slice operations have trivial implementations, so we do not show them here.

See Also

- The Server-Side main Function in C++
- C++ Mapping for Classes
- Transactional Evictor

Copyright 2016, ZeroC, Inc.

Adding an Evictor to the Java File System Server

On this page:

- The Server Main Program in Java
- The Persistent Servant Class Definitions in Java
- Implementing a Persistent FileI in Java
- Implementing a Persistent DirectoryI in Java
- Implementing NodeFactory in Java

The Server Main Program in Java

The server's `main` method is responsible for creating the evictor and initializing the root directory node. Many of the administrative duties, such as creating and destroying a communicator, are handled by the `Ice.Application` class. Our server `main` program has now become the following:

Java

```
import Filesystem.*;

public class Server extends Ice.Application
{
    public
    Server(String envName)
    {
        _envName = envName;
    }

    public int
    run(String[] args)
    {
        Ice.ObjectFactory factory = new NodeFactory();
        communicator().addObjectFactory(factory, PersistentFile.ice_staticId());
        communicator().addObjectFactory(factory, PersistentDirectory.ice_staticId()
);

        Ice.ObjectAdapter adapter = communicator().createObjectAdapter("EvictorFile
system");

        Freeze.Evictor evictor =
            Freeze.Util.createTransactionalEvictor(adapter, _envName, "evictorfs",
                                                   null, null, null, true);
        DirectoryI._evictor = evictor;
        FileI._evictor = evictor;

        adapter.addServantLocator(evictor, "");

        Ice.Identity rootId = new Ice.Identity();
        rootId.name = "RootDir";
        if(!evictor.hasObject(rootId))
        {
            PersistentDirectory root = new DirectoryI();
            root.nodeName = "/";
            root.nodes = new java.util.HashMap<java.lang.String, NodeDesc>();
            evictor.add(root, rootId);
        }
```


Copyright 2016, ZeroC, Inc.

```
        adapter.activate();

        communicator().waitForShutdown();

        return 0;
    }

    public static void
    main(String[] args)
    {
        Server app = new Server("db");
        int status = app.main("Server", args, "config.server");
        System.exit(status);
    }
```


Copyright 2016, ZeroC, Inc.

```
    private String _envName;
}
```

Let us examine the changes in detail. First, we define the class `Server` as a subclass of `Ice.Application`, and provide a constructor taking a string argument:

Java

```
    public
    Server(String envName)
    {
        _envName = envName;
    }
```

The string argument represents the name of the database environment, and is saved for later use in `run`.

One of the first tasks `run` performs is installing the Ice object factories for `PersistentFile` and `PersistentDirectory`. Although these classes are not exchanged via Slice operations, they are marshalled and unmarshalled in exactly the same way when saved to and loaded from the database, therefore factories are required. A single instance of `NodeFactory` is installed for both types:

Java

```
        Ice.ObjectFactory factory = new NodeFactory();
        communicator().addObjectFactory(factory, PersistentFile.ice_staticId());
        communicator().addObjectFactory(factory, PersistentDirectory.ice_staticId()
);
```

After creating the object adapter, the program initializes a transactional evictor by invoking `createTransactionalEvictor`. The third argument to `createTransactionalEvictor` is the name of the database, the fourth is null to indicate that we do not use facets, the fifth is null to indicate that we do not use a servant initializer, the sixth argument (`null`) indicates no indexes are in use, and the `true` argument requests that the database be created if it does not exist. The evictor is then added to the object adapter as a servant locator for the default category:

Java

```
        Freeze.Evictor evictor =
            Freeze.Util.createTransactionalEvictor(adapter, _envName, "evictorfs",
                                                   null, null, null, true);
        DirectoryI._evictor = evictor;
        FileI._evictor = evictor;

        adapter.addServantLocator(evictor, "");
```

Next, the program creates the root directory node if it is not already being managed by the evictor:

Copyright 2016, ZeroC, Inc.

Java

```
        Ice.Identity rootId = new Ice.Identity();
        rootId.name = "RootDir";
        if(!evictor.hasObject(rootId))
        {
            PersistentDirectory root = new DirectoryI();
            root.nodeName = "/";
            root.nodes = new java.util.HashMap<String, NodeDesc>();
            evictor.add(root, rootId);
        }
```

Finally, the `main` function instantiates the `Server` class, passing `db` as the name of the database environment:

Java

```
    public static void
    main(String[] args)
    {
        Server app = new Server("db");
        int status = app.main("Server", args, "config.server");
        System.exit(status);
    }
```

The Persistent Servant Class Definitions in Java

The servant classes must also be changed to incorporate the Freeze evictor. The `FileI` class now has a static state member `_evictor`:

Java

```
import Filesystem.*;

public final class FileI extends PersistentFile
{
    public
    FileI()
    {
        _destroyed = false;
    }

    // Slice operations...

    public static Freeze.Evictor _evictor;
    private boolean _destroyed;
}
```


Copyright 2016, ZeroC, Inc.

The `DirectoryI` class has undergone a similar transformation:

Java

```
import Filesystem.*;

public final class DirectoryI extends PersistentDirectory
{
    public
    DirectoryI()
    {
        _destroyed = false;
        nodes = new java.util.HashMap<String, NodeDesc>();
    }

    // Slice operations...

    public static Freeze.Evictor _evictor;
    private boolean _destroyed;
}
```

Implementing a Persistent `FileI` in Java

The `FileI` methods are mostly trivial, because the Freeze evictor handles persistence for us.

Java

Copyright 2016, ZeroC, Inc.

```
    public synchronized String
    name(Ice.Current current)
    {
        if (_destroyed) {
            throw new Ice.ObjectNotExistException(current.id, current.facet, curren
t.operation);
        }

        return nodeName;
    }

    public void
    destroy(Ice.Current current)
        throws PermissionDenied
    {
        synchronized(this) {
            if (_destroyed) {
                throw new Ice.ObjectNotExistException(current.id, current.facet, cu
rrent.operation);
            }
            _destroyed = true;
        }

        //
        // Because we use a transactional evictor,
        // these updates are guaranteed to be atomic.
        //
        parent.removeNode(nodeName);
        _evictor.remove(current.id);
    }

    public synchronized String[]
    read(Ice.Current current)
    {
        if (_destroyed) {
            throw new Ice.ObjectNotExistException(current.id, current.facet, curren
t.operation);
        }

        return (String[])text.clone();
    }

    public synchronized void
    write(String[] text, Ice.Current current)
        throws GenericError
    {
        if (_destroyed) {
            throw new Ice.ObjectNotExistException(current.id, current.facet, curren
t.operation);
        }

        this.text = text;
    }
```


Copyright 2016, ZeroC, Inc.

The code checks that the node has not been destroyed before acting on the invocation by updating or returning state. Note that `destroy` must update two separate nodes: as well as removing itself from the evictor, the node must also update the parent's node map. Because we are using a transactional evictor, the two updates are guaranteed to be atomic, so it is impossible to the leave the file system in an inconsistent state.

Implementing a Persistent `DirectoryI` in Java

The `DirectoryI` implementation requires more substantial changes. We begin our discussion with the `createDirectory` operation:

Java

```
    public synchronized DirectoryPrx
    createDirectory(String name, Ice.Current current)
        throws NameInUse
    {
        if (_destroyed) {
            throw new Ice.ObjectNotExistException(current.id, current.facet, curren
t.operation);
        }

        if (name.length() == 0 || nodes.containsKey(name)) {
            throw new NameInUse(name);
        }

        Ice.Identity id = current.adapter.getCommunicator().stringToIdentity(
            java.util.UUID.randomUUID().toString());
        PersistentDirectory dir = new DirectoryI();
        dir.nodeName = name;
        dir.parent = PersistentDirectoryPrxHelper.uncheckedCast(
            current.adapter.createProxy(current.id));
        DirectoryPrx proxy = DirectoryPrxHelper.uncheckedCast(_evictor.add(dir, id)
);

        NodeDesc nd = new NodeDesc();
        nd.name = name;
        nd.type = NodeType.DirType;
        nd.proxy = proxy;
        nodes.put(name, nd);

        return proxy;
    }
```

After validating the node name, the operation obtains a unique identity for the child directory, instantiates the servant, and registers it with the Freeze evictor. Finally, the operation creates a proxy for the child and adds the child to its node table.

The implementation of the `createFile` operation has the same structure as `createDirectory`:

Copyright 2016, ZeroC, Inc.

Java

```
    public synchronized FilePrx
    createFile(String name, Ice.Current current)
        throws NameInUse
    {
        if (_destroyed) {
            throw new Ice.ObjectNotExistException(current.id, current.facet, curren
t.operation);
        }

        if (name.length() == 0 || nodes.containsKey(name)) {
            throw new NameInUse(name);
        }

        Ice.Identity id = current.adapter.getCommunicator().stringToIdentity(
            java.util.UUID.randomUUID().toString());
        PersistentFile file = new FileI();
        file.nodeName = name;
        file.parent = PersistentDirectoryPrxHelper.uncheckedCast(
            current.adapter.createProxy(current.id));
        FilePrx proxy = FilePrxHelper.uncheckedCast(_evictor.add(file, id));

        NodeDesc nd = new NodeDesc();
        nd.name = name;
        nd.type = NodeType.FileType;
        nd.proxy = proxy;
        nodes.put(name, nd);

        return proxy;
    }
```

The remaining Slice operations have trivial implementations, so we do not show them here.

Implementing `NodeFactory` in Java

We use a single factory implementation for creating two types of Ice objects: `PersistentFile` and `PersistentDirectory`. These are the only two types that the Freeze evictor will be restoring from its database.

Copyright 2016, ZeroC, Inc.

Java

```
package Filesystem;

public class NodeFactory implements Ice.ObjectFactory
{
    public Ice.Object
    create(String type)
    {
        if (type.equals(PersistentFile.ice_staticId()))
            return new FileI();
        else if (type.equals(PersistentDirectory.ice_staticId()))
            return new DirectoryI();
        else {
            assert(false);
            return null;
        }
    }

    public void
    destroy()
    {
    }
}
```

See Also

- The Server-Side main Method in Java
- Java Mapping for Classes
- Transactional Evictor

Copyright 2016, ZeroC, Inc.

Freeze Maps

A Freeze map is a persistent, associative container in which the key type must be a legal dictionary key type and the value type can be any primitive or user-defined Slice type. For each pair of key and value types, the developer uses a code-generation tool to produce a language-specific class that conforms to the standard conventions for maps in that language. For example, in C++, the generated class resembles a `std::map`, and in Java it implements the `java.util.SortedMap` interface. Most of the logic for storing and retrieving state to and from the database is implemented in a Freeze base class. The generated map classes derive from this base class, so they contain little code and therefore are efficient in terms of code size.

You can only store data types that are defined in Slice in a Freeze map. Types without a Slice definition (that is, arbitrary C++ or Java types) cannot be stored because a Freeze map reuses the Ice-generated marshaling code to create the persistent representation of the data in the database. This is especially important to remember when defining a Slice class whose instances will be stored in a Freeze map; only the "public" (Slice-defined) data members will be stored, not the private state members of any derived implementation class.

Topics

- Freeze Map Concepts
- Using a Freeze Map in C++
- slice2freeze Command-Line Options
- Using a Freeze Map in Java
- slice2freezej Command-Line Options
- Using a Freeze Map in the File System Server

See Also

- Classes

Copyright 2016, ZeroC, Inc.

Freeze Map Concepts

On this page:

- Freeze Connections
- Using Transactions with Freeze Maps
  - Using Transactions with C++
  - Using Transactions with Java
- Iterating over a Freeze Map
- Recovering from Freeze Map Deadlocks
- Key Sorting for Freeze Maps
  - Key Sorting for Freeze Maps in C++
  - Key Sorting for Freeze Maps in Java
- Indexing a Freeze Map

Freeze Connections

In order to create a Freeze map object, you first need to obtain a Freeze `Connection` object by connecting to a database environment.

As illustrated in the following figure, a Freeze map is associated with a single connection and a single database file. Connection and map objects are not thread-safe: if you want to use a connection or any of its associated maps from multiple threads, you must serialize access to them. If your application requires concurrent access to the same database file (persistent map), you must create several connections and associated maps.

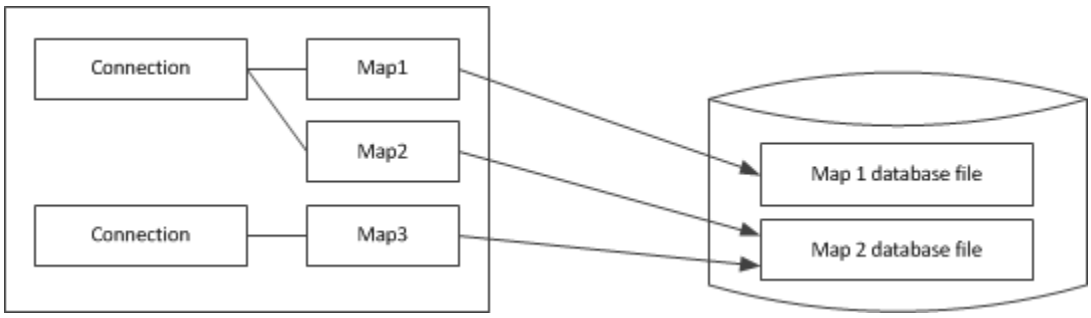

Freeze connections and maps.

Freeze connections provide operations that allow you to begin a transaction, access the current transaction, get the communicator associated with a connection, close a connection, and remove a map index. See the Slice API reference for more information on these operations.

Using Transactions with Freeze Maps

You may optionally use transactions with Freeze maps. Freeze transactions provide the usual ACID (atomicity, concurrency, isolation, durability) properties. For example, a transaction allows you to group several database updates in one atomic unit: either all or none of the updates within the transaction occur.

You start a transaction by calling `beginTransaction` on the `Connection` object. Once a connection has an associated transaction, all operations on the map objects associated with this connection use this transaction. Eventually, you end the transaction by calling `commit` or `rollback`: `commit` saves all your updates while `rollback` undoes them. The `currentTransaction` operation returns the transaction associated with a connection, if any; otherwise, it returns nil.

Copyright 2016, ZeroC, Inc.

Slice

```
module Freeze {

local interface Transaction {
    void commit();
    void rollback();
};

local interface Connection {
    Transaction beginTransaction();
    idempotent Transaction currentTransaction();
    // ...
};
};
```

If you do not use transactions, every non-iterator update is enclosed in its own internal transaction, and every read-write iterator has an associated internal transaction that is committed when the iterator is closed.

Using Transactions with C++

You must ensure that you either commit or roll back each transaction that you begin (otherwise, locks will be held by the database until they time out):

C++

```
ConnectionPtr connection = ...;

TransactionPtr tx = connection->beginTransaction();
try {

    // DB updates that might throw here...

    tx->commit();

    // More code that might throw here...

} catch (...) {
    try {
        tx->rollback();
    } catch (...) {
    }
    throw;
}
```

The outer try-catch blocks are necessary because, if the code encounters an exception, we must roll back any updates that were made. In turn, the attempt to roll back might throw itself, namely, if the code following the commit throws an exception (in which case the transaction cannot be rolled back because it is already committed).

Code such as this is difficult to maintain: for example, an early return statement can cause the transaction to be neither committed nor rolled back. The `TransactionHolder` class ensures that such errors cannot happen:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Freeze {
    class TransactionHolder {
    public:
        TransactionHolder(const ConnectionPtr&);
        ~TransactionHolder();

        void commit();
        void rollback();

    private:
        // Copy and assignment are forbidden.
        TransactionHolder(const TransactionHolder&);
        TransactionHolder& operator=(const TransactionHolder&);
    };
}
```

The constructor calls `beginTransaction` if the connection does not already have a transaction in progress, so instantiating the holder also starts a transaction. When the holder instance goes out of scope, its destructor calls `rollback` on the transaction and suppresses any exceptions that the rollback attempt might throw. This ensures that the transaction is rolled back if it was not previously committed or rolled back and ensures that an early return or an exception cannot cause the transaction to remain open:

C++

```
ConnectionPtr connection = ...;

{ // Open scope

    TransactionHolder tx(connection); // Begins transaction

    // DB updates that might throw here...

    tx.commit();

    // More code that might throw here...

} // Transaction rolled back here if not previously
  // committed or rolled back.
```

If you instantiate a `TransactionHolder` when a transaction is already in progress, it does nothing: the constructor notices that it could not begin a new transaction and turns `commit`, `rollback`, and the destructor into no-ops. For example, the nested `TransactionHolder` inst ance in the following code is benign and does nothing:

Copyright 2016, ZeroC, Inc.

C++

```
ConnectionPtr connection = ...;

{ // Open scope

    TransactionHolder tx(connection); // Begins transaction

    // DB updates that might throw here...

    { // Open nested scope

        TransactionHolder tx2(connection); // Does nothing

        // DB updates that might throw here...

        tx2.commit(); // Does nothing

        // More code that might throw here...

    } // Destructor of tx2 does nothing

    tx.commit();

    // More code that might throw here...

} // Transaction rolled back here if not previously
  // committed or rolled back.
```

Using Transactions with Java

You must ensure that you either commit or roll back each transaction that you begin (otherwise, locks will be held by the database until they time out):

Copyright 2016, ZeroC, Inc.

Java

```
Connection connection = ...;

Transaction tx = connection.beginTransaction();
try {

    // DB updates that might throw here...

    tx.commit();

    // More code that might throw here...

} catch (java.lang.RuntimeException ex) {
    try {
        tx.rollback();
    } catch (DatabaseException e) {
    }
    throw ex;
}
```

The catch handler ensures that the transaction is rolled back before re-throwing the exception. Note that the nested try-catch blocks are necessary: if the transaction committed successfully but the code following the commit throws an exception, the rollback attempt will fail therefore we need to suppress the corresponding `DatabaseException` that is raised in that case.

Also use caution with early `return` statements:

Copyright 2016, ZeroC, Inc.

Java

```
Connection connection = ...;

Transaction tx = connection.beginTransaction();
try {

    // DB updates that might throw here...

    if (error) {
        // ...
        return; // Oops, bad news!
    }

    // ...

    tx.commit();

    // More code that might throw here...

} catch (java.lang.RuntimeException ex) {
    try {
        tx.rollback();
    } catch (DatabaseException e) {
    }
    throw ex;
}
```

The early `return` statement in the preceding code causes the transaction to be neither committed nor rolled back. To deal with this situation, avoid early return statements or ensure that you either commit or roll back the transaction before returning. Alternatively, you can use a `finally` block to ensure that the transaction is rolled back:

Copyright 2016, ZeroC, Inc.

Java

```
Connection connection = ...;

try {

    Transaction tx = connection.beginTransaction();

    // DB updates that might throw here...

    if (error) {
        // ...
        return; // No problem, see finally block.
    }

    // ...

    tx.commit();

    // More code that might throw here...

} finally {
    if (connection.currentTransaction() != null)
        connection.currentTransaction().rollback();
}
```

Iterating over a Freeze Map

Iterators allow you to traverse the contents of a Freeze map. Iterators are implemented using Berkeley DB cursors and acquire locks on the underlying database page files. In C++, both read-only (`const_iterator`) and read-write iterators (`iterator`) are available. In Java, an iterator is read-write if it is obtained in the context of a transaction and read-only if it is obtained outside a transaction.

Locks held by an iterator are released when the iterator is closed (if you do not use transactions) or when the enclosing transaction ends. Releasing locks held by iterators is very important to let other threads access the database file through other connection and map objects. Occasionally, it is even necessary to release locks to avoid self-deadlock (waiting forever for a lock held by an iterator created by the same thread).

To improve ease of use and make self-deadlocks less likely, Freeze often closes iterators automatically. If you close a map or connection, associated iterators are closed. Similarly, when you start or end a transaction, Freeze closes all the iterators associated with the corresponding maps. If you do not use transactions, any write operation on a map (such as inserting a new element) automatically closes all iterators opened on the same map object, except for the current iterator when the write operation is performed through that iterator. In Java, Freeze also closes a read-only iterator when no more elements are available.

There is, however, one situation in C++ where an explicit iterator close is needed to avoid self-deadlock:

- you do not use transactions, and
- you have an open iterator that was used to update a map (it holds a write lock), and
- in the same thread, you read that map.

Read operations in C++ never close iterators automatically: you need to either use transactions or explicitly close the iterator that holds the write lock. This is not an issue in Java because you cannot use an iterator to update a map outside of a transaction.

Recovering from Freeze Map Deadlocks

If you use multiple threads to access a database file, Berkeley DB may acquire locks in conflicting orders (on behalf of different transactions or iterators). For example, an iterator could have a read-lock on page P1 and attempt to acquire a write-lock on page P2, while another iterator (on a different map object associated with the same database file) could have a read-lock on P2 and attempt to acquire a write-lock on P1.

Copyright 2016, ZeroC, Inc.

When this occurs, Berkeley DB detects a deadlock and resolves it by returning a "deadlock" error to one or more threads. For all non-iterator operations performed outside any transaction, such as an insertion into a map, Freeze catches such errors and automatically retries the operation until it succeeds. (In that case, the most-recently acquired lock is released before retrying.) For other operations, Freeze reports this deadlock by raising `Freeze::DeadlockException`. In that case, the associated transaction or iterator is also automatically rolled back or closed. A properly written application must expect to catch deadlock exceptions and retry the transaction or iteration.

Key Sorting for Freeze Maps

Keys in Freeze maps and indexes are always sorted. By default, Freeze sorts keys according to their Ice-encoded binary representation; this is very efficient but the resulting order is rarely meaningful for the application. Starting with Ice 3.0, Freeze offers the ability to specify your own comparator objects so that you can customize the traversal order of your maps. Note however that the comparator of a Freeze map should remain the same throughout the life of the map. Berkeley DB stores records according to the key order provided by this comparator; switching to another comparator will cause undefined behavior.

Key Sorting for Freeze Maps in C++

In C++, you specify the name of your comparator objects during code generation. The generated map provides the standard features of `std::map`, so that iterators return entries according to the order you have defined for the main key with your comparator object. The `lower_bound`, `upper_bound`, and `equal_range` functions provide range-searches (see the definition of these functions on `std::map`).

Apart from these standard features, the generated map provides additional functions and methods to perform range searches using secondary keys. The additional functions are `lowerBoundFor`*`Member`*, `upperBoundFor`*`Member`*, and `equalRangeFor`*`Member`*, where *`Member`* is the name of the secondary-key member. These functions return regular iterators on the Freeze map.

Key Sorting for Freeze Maps in Java

In Java, you supply comparator objects (instances of the standard Java interface `java.util.Comparator`) at run time when instantiating the generated map class. The map constructor accepts a comparator for the main key and optionally a collection of comparators for secondary keys. The map also provides a number of methods for performing range searches on the main key and on secondary keys.

Indexing a Freeze Map

Freeze maps support efficient reverse lookups: if you define an index when you generate your map (with `slice2freeze` or `slice2freezej`), the generated code provides additional methods for performing reverse lookups. If your value type is a structure or a class, you can also index on a member of the value, and several such indices can be associated with the same Freeze map.

Indexed searches are easy to use and very efficient. However, be aware that an index also adds significant write overhead: with Berkeley DB, every update triggers a read from the database to get the old index entry and, if necessary, replace it.

If you add an index to an existing map, Freeze automatically populates the index the next time you open the map. Freeze populates the index by instantiating each map entry, so it is important that you register the object factories for any class types in your map before you open the map.

Note that the index key comparator of a Freeze map index should remain the same throughout the life of the index. Berkeley DB stores records according to the key order provided by this comparator; switching to another comparator will cause undefined behavior.

See Also

- Using a Freeze Map in C++
- Using a Freeze Map in Java

Copyright 2016, ZeroC, Inc.

Using a Freeze Map in C++

This page describes the C++ code generator and demonstrates how to use a Freeze map in a C++ program.

On this page:

- Generating a Simple Map for C++
- The Freeze Map Class in C++
- Using Iterators with Freeze Maps in C++
- Sample Freeze Map Program in C++

Generating a Simple Map for C++

We can use slice2freeze to generate a simple Freeze map:

```
$ slice2freeze --dict StringIntMap,string,int StringIntMap
```

This command directs the compiler to create a map named `StringIntMap`, with the Slice key type `string` and the Slice value type `int`. The final argument is the base name for the output files, to which the compiler appends the `.h` and `.cpp` suffixes. As a result, this command produces two C++ source files, `StringIntMap.h` and `StringIntMap.cpp`.

The Freeze Map Class in C++

If you examine the contents of the header file created by the example in the previous section, you will discover that a Freeze map is a subclass of the template class `Freeze::Map`:

C++

```
// StringIntMap.h

class StringIntMap : public Freeze::Map< std::string, Ice::Int, ... > {
public:
    StringIntMap(const Freeze::ConnectionPtr& connection,
                 const std::string& dbName,
                 bool createDb = true,
                 const Freeze::IceEncodingCompare& compare = ...);

    template <class _InputIterator>
    StringIntMap(const Freeze::ConnectionPtr& connection,
                 const std::string& dbName,
                 bool createDb,
                 _InputIterator first, _InputIterator last,
                 const Freeze::IceEncodingCompare& compare = ...);

    ...
};
```

The `Freeze::Map` template class closely resembles the STL container class `std::map`, as shown in the following class definition:

C++

```
namespace Freeze {
```


Copyright 2016, ZeroC, Inc.

```
template<...> class Map {
public:
    typedef ... value_type;
    typedef ... iterator;
    typedef ... const_iterator;

    typedef size_t size_type;
    typedef ptrdiff_t difference_type;

    static void recreate(const Freeze::ConnectionPtr& connection,
                         const std::string& dbName,
                         const Freeze::IceEncodingCompare& compare = ...);

    bool operator==(const Map& rhs) const;
    bool operator!=(const Map& rhs) const;

    void swap(Map& rhs);

    iterator begin();
    const_iterator begin() const;

    iterator end();
    const_iterator end() const;

    bool empty() const;
    size_type size() const;
    size_type max_size() const;

    iterator insert(iterator /*position*/, const value_type& elem);

    std::pair<iterator, bool> insert(const value_type& elem);

    template <typename InputIterator>
    void insert(InputIterator first, InputIterator last);

    void put(const value_type& elem);

    template <typename InputIterator>
    void put(InputIterator first, InputIterator last);

    void erase(iterator position);
    size_type erase(const key_type& key);
    void erase(iterator first, iterator last);

    void clear();

    void destroy(); // Non-standard.

    iterator find(const key_type& key);
    const_iterator find(const key_type& key) const;

    size_type count(const key_type& key) const;

    iterator lower_bound(const key_type& key);
    const_iterator lower_bound(const key_type& key) const;
    iterator upper_bound(const key_type& key);
```


Copyright 2016, ZeroC, Inc.

```
const_iterator upper_bound(const key_type& key) const;

std::pair<iterator, iterator>
equal_range(const key_type& key);

std::pair<const_iterator, const_iterator>
equal_range(const key_type& key) const;

const Ice::CommunicatorPtr&
communicator() const;

...
```


Copyright 2016, ZeroC, Inc.

```
};
}
```

The semantics of the `Freeze::Map` methods are identical to those of `std::map` unless otherwise noted. In particular, the overloaded `insert` method shown below ignores the `position` argument:

C++

```
iterator insert(iterator /*position*/, const value_type& elem);
```

A Freeze map class supports only those methods shown above; other features of `std::map`, such as allocators and overloaded array operators, are not available.

Non-standard methods that are specific to Freeze maps are discussed below:

- Constructors
  The following overloaded constructors are provided:

  C++

  ```
  Map(const Freeze::ConnectionPtr& connection,
      const std::string& dbName,
      bool createDb = true,
      const Freeze::IceEncodingCompare& compare = ...);

  template<class _InputIterator>
  Map(const Freeze::ConnectionPtr& connection,
      const std::string& dbName,
      bool createDb,
      _InputIterator first, _InputIterator last,
      const Freeze::IceEncodingCompare& compare = ...);
  ```

  The first constructor accepts a connection, the database name, a flag indicating whether to create the database if it does not exist, and an object used to compare keys. The second constructor accepts all of the parameters of the first, with the addition of iterators from which elements are copied into the map.

  Note that a database can only contain the persistent state of one map type. Any attempt to instantiate maps of different types on the same database results in undefined behavior.

- Map copy
  The `recreate` function copies an existing database:

  C++

  ```
  static void recreate(const Freeze::ConnectionPtr& connection,
                       const std::string& dbName,
                       const Freeze::IceEncodingCompare& compare = ...)
  ```

  The `dbName` parameter specifies an existing database name. The copy has the name <*dbName*>`.old-`<*uuid*>. For example, if the

Copyright 2016, ZeroC, Inc.

database name is `MyDB`, the copy might be named `MyDB.old-edefd55a-e66a-478d-a77b-f6d53292b873`. (Obviously, a different UUID is used each time you recreate a database).

- `destroy`
  This method deletes the database from its environment and from the Freeze catalog. If a transaction is not currently open, the method creates its own transaction in which to perform this task.
- `communicator`
  This method returns the communicator with which the map's connection is associated.

Using Iterators with Freeze Maps in C++

A Freeze map's iterator works like its counterpart in `std::map`. The iterator class supports one convenient (but nonstandard) method:

C++

```
void set(const mapped_type& value)
```

Using this method, a program can replace the value at the iterator's current position.

Sample Freeze Map Program in C++

The program below demonstrates how to use a `StringIntMap` to store `<string, int>` pairs in a database. You will notice that there are no explicit `read` or `write` operations called by the program; instead, simply using the map has the side effect of accessing the database.

C++

Copyright 2016, ZeroC, Inc.

```
#include <Freeze/Freeze.h>
#include <StringIntMap.h>

int
main(int argc, char* argv[])
{
    // Initialize the Communicator.
    //
    Ice::CommunicatorPtr communicator = Ice::initialize(argc, argv);

    // Create a Freeze database connection.
    //
    Freeze::ConnectionPtr connection =
Freeze::createConnection(communicator, "db");

    // Instantiate the map.
    //
    StringIntMap map(connection, "simple");

    // Clear the map.
    //
    map.clear();

    Ice::Int i;
    StringIntMap::iterator p;

    // Populate the map.
    //
    for (i = 0; i < 26; i++) {
        std::string key(1, 'a' + i);
        map.insert(make_pair(key, i));
    }

    // Iterate over the map and change the values.
    //
    for (p = map.begin(); p != map.end(); ++p)
        p.set(p->second + 1);

    // Find and erase the last element.
    //
    p = map.find("z");
    assert(p != map.end());
    map.erase(p);

    // Clean up.
    //
    connection->close();
    communicator->destroy();

    return 0;
}
```

Prior to instantiating a Freeze map, the application must connect to a Berkeley DB database environment:

Copyright 2016, ZeroC, Inc.

C++

```
Freeze::ConnectionPtr connection = Freeze::createConnection(communicator, "db");
```

The second argument is the name of a Berkeley DB database environment; by default, this is also the file system directory in which Berkeley DB creates all database and administrative files. Note that properties with the prefix `Freeze.DbEnv` can modify a number of environment settings, including the file system directory. For the preceding example, you could change the directory to `FreezeDir` by setting the property `Freeze.DbEnv.db.DbHome` to `FreezeDir`.

Next, the code instantiates the `StringIntMap` on the connection. The constructor's second argument supplies the name of the database file, which by default is created if it does not exist:

C++

```
StringIntMap map(connection, "simple");
```

After instantiating the map, we clear it to make sure it is empty in case the program is run more than once:

C++

```
map.clear();
```

Next, we populate the map using a single-character string as the key:

C++

```
for (i = 0; i < 26; i++) {
    std::string key(1, 'a' + i);
    map.insert(make_pair(key, i));
}
```

Iterating over the map will look familiar to `std::map` users. However, to modify a value at the iterator's current position, we use the nonstandard `set` method:

C++

```
for (p = map.begin(); p != map.end(); ++p)
    p.set(p->second + 1);
```

Next, the program obtains an iterator positioned at the element with key `z`, and erases it:

Copyright 2016, ZeroC, Inc.

C++

```
p = map.find("z");
assert(p != map.end());
map.erase(p);
```

Finally, the program closes the database connection:

C++

```
connection->close();
```

It is not necessary to explicitly close the database connection, but we demonstrate it here for the sake of completeness.

See Also

- Using the Slice Compilers
- slice2freeze Command-Line Options
- Freeze Map Concepts

Copyright 2016, ZeroC, Inc.

slice2freeze Command-Line Options

The Slice-to-Freeze compiler, `slice2freeze`, generates C++ classes for Freeze maps and Freeze Evictor indices.

`slice2freeze`'s first command-line parameter is the base name for the generated C++ files, and not a Slice file like with the other Slice compilers:

```
slice2freeze [options] BaseName [SliceFile1 SliceFile2 ...]
```

The compiler offers the following command-line options in addition to the standard options:

--header-ext EXT

Changes the file extension for the generated header files from the default h to the extension specified by *EXT*.

--source-ext EXT

Changes the file extension for the generated source files from the default `cpp` to the extension specified by *EXT*.

--add-header HDR[,GUARD]

This option adds an include directive for the specified header at the beginning of the generated source file (preceding any other include directives). If *GUARD* is specified, the include directive is protected by the specified guard.

For example:

```
--add-header precompiled.h,_PRECOMPILED_H__
```

results in the following directives at the beginning of the generated source file:

C++

```
#ifndef __PRECOMPILED_H__
#define __PRECOMPILED_H__
#include <precompiled.h>
#endif
```

As this example demonstrates, the `--add-header` option is useful mainly for integrating the generated code with a compiler's precompiled header mechanism.

This option can be repeated to create include directives for several files.

--include-dir DIR

Modifies `#include` directives in source files to prepend the path name of each header file with the directory *DIR*. The discussion of `slice2cpp` provides more information.

--dll-export SYMBOL

Use *SYMBOL* to control DLL exports or imports. See the `slice2cpp` description for details.

--dict NAME,KEY,VALUE[,sort[,COMPARE]]

Generate a Freeze map class named *NAME* using *KEY* as key and *VALUE* as value. This option may be specified multiple times to generate

Copyright 2016, ZeroC, Inc.

several Freeze maps. *NAME* may be a scoped C++ name, such as `Demo::Struct1ObjectMap`. *KEY* and *VALUE* represent Slice types and therefore must use Slice syntax, such as `bool` or `Ice::Identity`. The type identified by *KEY* must be a legal dictionary key type. By default, keys are sorted using their binary Ice-encoded representation. Include `sort` to sort with the *COMPARE* functor class. If *COMPARE* is not specified, the default value is `std::less<`*KEY*`>`.

--dict-index MAP[,MEMBER] [,case-sensitive|case-insensitive][,sort[,COMPARE]]

Add an index to the Freeze map named *MAP*. If *MEMBER* is specified, the map value type must be a structure or a class, and *MEMBER* must be a member of this structure or class. Otherwise, the entire value is indexed. When the indexed member (or entire value) is a string, the index can be case-sensitive (default) or case-insensitive. An index adds additional member functions to the generated C++ map:

- `iterator findBy`*MEMBER*`(`*MEMBER_TYPE*`, bool = true);`
- `const_iterator findBy`*MEMBER*`(`*MEMBER_TYPE*`,bool = true) const;`
- `iterator beginFor`*MEMBER*`();`
- `const_iterator begin_For`*MEMBER*`() const;`
- `iterator endFor`*MEMBER*`();`
- `const_iterator endFor`*MEMBER*`() const;`
- `iterator lowerBoundFor`*MEMBER*`(`*MEMBER_TYPE*`);`
- `const_iterator lowerBoundFor`*MEMBER*`(`*MEMBER_TYPE*`) const;`
- `iterator upperBoundFor`*MEMBER*`(`*MEMBER_TYPE*`);`
- `const_iterator upperBoundFor`*MEMBER*`(`*MEMBER_TYPE*`) const;`
- `std::pair<iterator, iterator> equalRangeFor`*MEMBER*`(`*MEMBER_TYPE*`);`
- `std::pair<const_iterator, const_iterator> equalRangeFor`*MEMBER*`(`*MEMBER_TYPE*`) const;`
- `int` *MEMBER*`Count(`*MEMBER_TYPE*`) const;`

When *MEMBER* is not specified, these functions are `findByValue` (const and non-const), `lowerBoundForValue` (const and non-const), `valueCount`, and so on. When *MEMBER* is specified, its first letter is capitalized in the `findBy` function name. *MEMBER_TYPE* corresponds to an in-parameter of the type of *MEMBER* (or the type of the value when *MEMBER* is not specified). For example, if *MEMBER* is a string, *MEMBER_TYPE* is a `const std::string&`.

By default, keys are sorted using their binary Ice-encoded representation. Include `sort` to sort with the *COMPARE* functor class. If *COMPARE* is not specified, the default value is `std::less<`*MEMBER_TYPE*`>`.

`findBy`*MEMBER* returns an iterator to the first element in the Freeze map that matches the given index value. It returns `end()` if there is no match. When the second parameter is true (the default), the returned iterator provides only the elements with an exact match (and then skips to `end()`). Otherwise, the returned iterator sets a starting position and then provides all elements until the end of the map, sorted according to the index comparator.

`lowerBoundFor`*MEMBER* returns an iterator to the first element in the Freeze map whose index value is not less than the given index value. It returns `end()` if there is no such element. The returned iterator provides all elements until the end of the map, sorted according to the index comparator.

`upperBoundFor`*MEMBER* returns an iterator to the first element in the Freeze map whose index value is greater than the given index value. It returns `end()` if there is no such element. The returned iterator provides all elements until the end of the map, sorted according to the index comparator.

`beginFor`*MEMBER* returns an iterator to the first element in the map.

`endFor`*MEMBER* returns an iterator to the last element in the map.

`equalRangeFor`*MEMBER* returns a range (pair of iterators) of all the elements whose index value matches the given index value. This function is similar to `findBy`MEMBER (see above).

MEMBER`Count` returns the number of elements in the Freeze map whose index value matches the given index value.

Please note that index-derived iterators do not allow you to set new values in the underlying map.

--index CLASS,TYPE,MEMBER [,case-sensitive|case-insensitive]

Generate an index class for a Freeze evictor. *CLASS* is the name of the class to be generated. *TYPE* denotes the type of class to be indexed (objects of different classes are not included in this index). *MEMBER* is the name of the data member in *TYPE* to index. When *MEMBER* has type `string`, it is possible to specify whether the index is case-sensitive or not. The default is case-sensitive.

See Also

- Using the Slice Compilers
- slice2cpp Command-Line Options
- Freeze Map Concepts

Copyright 2016, ZeroC, Inc.

Using a Freeze Map in Java

This page describes how to generate and use a Freeze map in a Java program.

On this page:

- Generating a Simple Map for Java
- The Freeze Map Class in Java
- Why Comparators are Important
- Using Iterators with Freeze Maps in Java
- Generating Indices for Freeze Maps in Java
- Sample Freeze Map Program in Java

Generating a Simple Map for Java

We can use slice2freezej to generate a simple Freeze map:

```
$ slice2freezej --dict StringIntMap,string,int
```

This command directs the compiler to create a map named `StringIntMap`, with the Slice key type `string` and the Slice value type `int`. The compiler produces one Java source file: `StringIntMap.java`.

The Freeze Map Class in Java

The class generated by `slice2freezej` implements the `Freeze.Map` interface, as shown below:

Java

```
package Freeze;

public interface Map<K, V> extends NavigableMap<K, V>
{
    void fastPut(K key, V value);
    void close();
    int closeAllIterators();
    void destroy();

    public interface EntryIterator<T> extends java.util.Iterator<T>
    {
        void close();
        void destroy(); // an alias for close
    }
}
```

The `Map` interface implements standard Java interfaces and provides nonstandard methods that improve efficiency and support database-oriented features. `Map` defines the following methods:

- `fastPut`
  Inserts a new key-value pair. This method is more efficient than the standard `put` method because it avoids the overhead of reading and decoding the previous value associated with the key (if any).
- `close`
  Closes the database associated with this map along with all open iterators. A map must be closed when it is no longer needed, either by closing the map directly or by closing the Freeze `Connection` object with which this map is associated.
- `closeAllIterators`
  Closes all open iterators and returns the number of iterators that were closed. We discuss iterators in more detail in the next section.

Copyright 2016, ZeroC, Inc.

- `destroy`
  Removes the database associated with this map along with any indices.

`Map` inherits much of its functionality from the `Freeze.NavigableMap` interface, which derives from the standard Java interface `java.util.SortedMap` and also supports a subset of the `java.util.NavigableMap` interface:

Java

```
package Freeze;

public interface NavigableMap<K, V> extends java.util.SortedMap<K, V>
{
    java.util.Map.Entry<K, V> firstEntry();
    java.util.Map.Entry<K, V> lastEntry();

    java.util.Map.Entry<K, V> ceilingEntry(K key);
    java.util.Map.Entry<K, V> floorEntry(K key);
    java.util.Map.Entry<K, V> higherEntry(K key);
    java.util.Map.Entry<K, V> lowerEntry(K key);

    K ceilingKey(K key);
    K floorKey(K key);
    K higherKey(K key);
    K lowerKey(K key);

    java.util.Set<K> descendingKeySet();
    NavigableMap<K, V> descendingMap();

    NavigableMap<K, V> headMap(K toKey, boolean inclusive);
    NavigableMap<K, V> tailMap(K fromKey, boolean inclusive);
    NavigableMap<K, V> subMap(K fromKey, boolean fromInclusive,
                              K toKey, boolean toInclusive);

    java.util.Map.Entry<K, V> pollFirstEntry();
    java.util.Map.Entry<K, V> pollLastEntry();

    boolean fastRemove(K key);
}
```

The `NavigableMap` interface provides a number of useful methods:

- `firstEntry`
  `lastEntry`
  Returns the first and last key-value pair, respectively.
- `ceilingEntry`
  Returns the key-value pair associated with the least key greater than or equal to the given key, or null if there is no such key.
- `floorEntry`
  Returns the key-value pair associated with the greatest key less than or equal to the given key, or null if there is no such key.
- `higherEntry`
  Returns the key-value pair associated with the least key greater than the given key, or null if there is no such key.
- `lowerEntry`
  Returns the key-value pair associated with the greatest key less than the given key, or null if there is no such key.
- `ceilingKey`
  `floorKey`

Copyright 2016, ZeroC, Inc.

`higherKey`
`lowerKey`
These methods have the same semantics as those described above, except they return only the key portion of the matching key-value pair or null if there is no such key.

- `descendingKeySet`
  Returns a set representing a reverse-order view of the keys in this map.
- `descendingMap`
  Returns a reverse-order view of the entries in this map.
- `headMap`
  Returns a view of the portion of this map whose keys are less than (or equal to, if inclusive is true) the given key.
- `tailMap`
  Returns a view of the portion of this map whose keys are greater than (or equal to, if inclusive is true) the given key.
- `subMap`
  Returns a view of the portion of this map whose keys are within the given range.
- `pollFirstEntry`
  `pollLastEntry`
  Removes and returns the first and last key-value pair, respectively.
- `fastRemove`
  Removes an existing key-value pair. As for `fastPut`, this method is a more efficient alternative to the standard `remove` method that returns true if a key-value pair was removed, or false if no match was found.

You must supply a comparator object when constructing the map in order to use many of these methods.

Note that `NavigableMap` also inherits overloaded methods named `headMap`, `tailMap`, and `subMap` from the `SortedMap` interface. These methods have the same semantics as the ones defined in `NavigableMap` but they omit the boolean arguments (refer to the JDK documentation for complete details). Although these methods are declared as returning a `SortedMap`, the actual type of the returned object is a `NavigableMap` that you can downcast if necessary.

There are some limitations in the sub maps returned by the `headMap`, `tailMap` and `subMap` methods:

- A new entry in the Freeze map cannot be added via a sub map, therefore calling `put` raises `UnsupportedOperationException`.
- An existing entry in the Freeze map cannot be removed via a sub map or iterator for a secondary key.

Now let us examine the contents of the source file created by the example in the previous section:

Copyright 2016, ZeroC, Inc.

Java

```
public class StringIntMap extends ...
    // implements Freeze.Map<String, Integer>
{
    public StringIntMap(
        Freeze.Connection connection,
        String dbName,
        boolean createDb,
        java.util.Comparator<String> comparator);

    public StringIntMap(
        Freeze.Connection connection,
        String dbName,
        boolean createDb);

    public StringIntMap(
        Freeze.Connection connection,
        String dbName);
}
```

`StringIntMap` derives from an internal Freeze base class that implements the interface `Freeze.Map<String, Integer>`. The generated class defines several overloaded constructors whose arguments are described below:

- `connection`
  The Freeze connection object.
- `dbName`
  The name of the database in which to store this map's persistent state. Note that a database can only contain the persistent state of one map type. Any attempt to instantiate maps of different types on the same database results in undefined behavior.
- `createDb`
  A flag indicating whether the map should create the database if it does not already exist. If this argument is not specified, the default value is `true`.
- `comparator`
  An object used to compare the map's keys. If this argument is not specified, the default behavior compares the encoded form of the keys.

Why Comparators are Important

The constructor of a Freeze map optionally accepts a comparator object for the primary key and, if any indices are generated, a second object that supplies comparators for each of the index keys. If you do not supply a comparator, Freeze simply compares the encoded form of the keys. This default behavior is acceptable when comparing keys for equality, but using the encoded form cannot work reliably when comparing keys for ordering purposes.

For example, many of the methods in `NavigableMap` perform greater-than or less-than comparisons on keys, including `ceilingEntry`, `headMap`, and `tailMapFor`*`MEMBER`*. All of these methods raise `UnsupportedOperationException` if you failed to supply a corresponding comparator when constructing the map. (The same applies to `NavigableMap` objects created for secondary keys.) In fact, the only `NavigableMap` methods that do *not* require a comparator are `firstEntry`, `lastEntry`, `pollFirstEntry`, `pollLastEntry`, and `fastRemove`.

As you can see, the functionality of a Freeze map is quite limited if no comparators are configured, therefore we recommend using comparators at all times.

Using Iterators with Freeze Maps in Java

You can iterate over a Freeze map just as you can with any container that implements the `java.util.Map` interface. For example, the

Copyright 2016, ZeroC, Inc.

code below displays the key and value of each element:

Java

```
StringIntMap m = new StringIntMap(...);
java.util.Iterator<java.util.Map.Entry<String, Integer>> i =
m.entrySet().iterator();
while (i.hasNext()) {
    java.util.Map.Entry<String, Integer> e = i.next();
    System.out.println("Key: " + e.getKey());
    System.out.println("Value: " + e.getValue());
}
```

Generally speaking, a program should close an iterator when it is no longer necessary. (An iterator that is garbage collected without being closed emits a warning message.) However, an explicit close was not necessary in the preceding example because Freeze automatically closes a read-only iterator when it reaches the last element (a read-only iterator is one that is opened outside of any transaction). If instead our program had stopped using the iterator prior to reaching the last element, an explicit close would have been necessary:

Java

```
StringIntMap m = new StringIntMap(...);
java.util.Iterator<java.util.Map.Entry<String, Integer>> i =
m.entrySet().iterator();
while (i.hasNext()) {
    java.util.Map.Entry<String, Integer> e = i.next();
    System.out.println("Key: " + e.getKey());
    System.out.println("Value: " + e.getValue());
    if (e.getValue().intValue() == 5)
        break;
}
((Freeze.Map.EntryIterator)i).close();
```

Closing the iterator requires downcasting it to a Freeze-specific interface named `Freeze.Map.EntryIterator`. The definition of this interface was shown in the previous section.

Freeze maps also support the enhanced `for` loop functionality. Here is a simpler way to write our original program:

Java

```
StringIntMap m = new StringIntMap(...);
for (java.util.Map.Entry<String, Integer> e : m.entrySet()) {
    System.out.println("Key: " + e.getKey());
    System.out.println("Value: " + e.getValue());
}
```

As in the first example, Freeze automatically closes the iterator when no more elements are available. Although the enhanced `for` loop is convenient, it is not appropriate for all situations because the loop hides its iterator and therefore prevents the program from accessing the iterator in order to close it. In this case, you can use the traditional `while` loop instead of the `for` loop, or you can invoke `closeAllIterators` on the map as shown below:

Copyright 2016, ZeroC, Inc.

Java

```
StringIntMap m = new StringIntMap(...);
for (java.util.Map.Entry<String, Integer> e : m.entrySet()) {
    System.out.println("Key: " + e.getKey());
    System.out.println("Value: " + e.getValue());
    if (e.getValue().intValue() == 5)
        break;
}
int num = m.closeAllIterators();
assert(num <= 1); // The iterator may already be closed.
```

The `closeAllIterators` method returns an integer representing the number of iterators that were actually closed. This value can be useful for diagnostic purposes, such as to assert that a program is correctly closing its iterators.

Generating Indices for Freeze Maps in Java

Using the `--dict-index` option to define an index for a secondary key causes `slice2freezej` to generate the following additional code in a Freeze map:

- A static nested class named `IndexComparators`, which allows you to supply a custom comparator object for each index in the map.
- An overloading of the map constructor that accepts an instance of `IndexComparators`.
- An overloading of the `recreate` method that accepts an instance of `IndexComparators`.
- Searching, counting, and range-searching methods for finding key-value pairs using the secondary key.

We discuss each of these additions in more detail below. In this discussion, *`MEMBER`* refers to the optional argument of the `--dict-index` option, and *`MEMBER_TYPE`* refers to the type of that member. As explained earlier, if *`MEMBER`* is not specified, `slice2freezej` creates an index for the value type of the map. The sample code presented in this section assumes we have generated a Freeze map using the following command:

Java

```
$ slice2freezej --dict StringIntMap,string,int --dict-index StringIntMap
```

By default, index keys are sorted using their binary Ice-encoded representation. This is an efficient sorting scheme but does not necessarily provide a meaningful traversal order for applications. You can choose a different order by providing an instance of the `IndexComparators` class to the map constructor. This class has a public data member holding a comparator (an instance of `java.util.Comparator<`*`MEMBER_TYPE`*`>`) for each index in the map. The class also provides an empty constructor as well as a convenience constructor that allows you to instantiate and initialize the object all at once. The name of each data member is *`MEMBER`*`Comparator`. If *`MEMBER`* is not specified, the `IndexComparators` class has a single data member named `valueComparator`.

Much of the functionality offered by a map index requires that you provide a custom comparator.

Here is the definition of `IndexComparators` for `StringIntMap`:

Copyright 2016, ZeroC, Inc.

Java

```
public class StringIntMap ... {
    public static class IndexComparators {
        public IndexComparators() {}

        public IndexComparators(java.util.Comparator<Integer> valueComparator);

        public java.util.Comparator<Integer> valueComparator;
    }

    ...
}
```

To instantiate a Freeze map using your custom comparators, you must use the overloaded constructor that accepts the `IndexComparators` object. For our `StringIntMap`, this constructor has the following definition:

Java

```
public class StringIntMap ... {
    public StringIntMap(
        Freeze.Connection connection,
        String dbName,
        boolean createDb,
        java.util.Comparator<String> comparator,
        IndexComparators indexComparators);

    ...
}
```

Now we can instantiate our `StringIntMap` as follows:

Java

```
java.util.Comparator<String> myMainKeyComparator = ...;
StringIntMap.IndexComparators indexComparators =
new StringIntMap.IndexComparators();
indexComparators.valueComparator = ...;
StringIntMap m = new StringIntMap(connection, "stringIntMap", true,
                                  myMainKeyComparator, indexComparators);
```

If you later need to change the index configuration of a Freeze map, you can use one of the `recreate` methods to update the database. Here are the definitions from `StringIntMap`:

Copyright 2016, ZeroC, Inc.

Java

```
public class StringIntMap ... {
    public static void recreate(
        Freeze.Connection connection,
        String dbName,
        java.util.Comparator<String> comparator);

    public static void recreate(
        Freeze.Connection connection,
        String dbName,
        java.util.Comparator<String> comparator,
        IndexComparators indexComparators);

    ...
}
```

The first overloading is generated for every map, whereas the second overloading is only generated when the map has at least one index. As its name implies, the `recreate` method creates a new copy of the database. More specifically, the method removes any existing indices, copies every key-value pair to a temporary database, and finally replaces the old database with the new one. As a side-effect, this process also populates any remaining indices. The first overloading of `recreate` is useful when you have regenerated the map to remove the last index and wish to clean up the map's database state.

`slice2freezej` also generates a number of index-specific methods. The names of these methods incorporate the member name (*`MEMBER`*), or use `value` if *`MEMBER`* is not specified. In each method name, the value of *`MEMBER`* is used unchanged if it appears at the beginning of the method's name. Otherwise, if *`MEMBER`* is used elsewhere in the method name, its first letter is capitalized. The index methods are described below:

- `public Freeze.Map.EntryIterator<Map.Entry<K, V>>`
  `findBy`*`MEMBER`*`(`*`MEMBER_TYPE`*` index)`

  `public Freeze.Map.EntryIterator<Map.Entry<K, V>>`
  `findBy`*`MEMBER`*`(`*`MEMBER_TYPE`*` index, boolean onlyDups)`
  Returns an iterator over elements of the Freeze map starting with an element with whose index value matches the given index value. If there is no such element, the returned iterator is empty (`hasNext` always returns false). When the second parameter is true (or is not provided), the returned iterator provides only "duplicate" elements, that is, elements with the very same index value. Otherwise, the iterator sets a starting position in the map, and then provides elements until the end of the map, sorted according to the index comparator. Any attempt to modify the map via this iterator results in an `UnsupportedOperationException`.

- `public int `*`MEMBER`*`Count(`*`MEMBER_TYPE`*` index)`
  Returns the number of elements in the Freeze map whose index value matches the given index value.

- `public NavigableMap<`*`MEMBER_TYPE`*`, Set<Map.Entry<K, V>>>`
  `headMapFor`*`MEMBER`*`(`*`MEMBER_TYPE`*` to, boolean inclusive)`

  `public NavigableMap<`*`MEMBER_TYPE`*`, Set<Map.Entry<K, V>>>`
  `headMapFor`*`MEMBER`*`(`*`MEMBER_TYPE`*` to)`
  Returns a view of the portion of the Freeze map whose keys are less than (or equal to, if `inclusive` is true) the given key. If `inclusive` is not specified, the method behaves as if `inclusive` is false.

- `public NavigableMap<`*`MEMBER_TYPE`*`, Set<Map.Entry<K, V>>>`
  `tailMapFor`*`MEMBER`*`(`*`MEMBER_TYPE`*` from, boolean inclusive)`

- public NavigableMap<*`MEMBER_TYPE`*, Set<Map.Entry<K, V>>>}}
  `tailMapFor`*`MEMBER`*`(`*`MEMBER_TYPE`*` from)`
  Returns a view of the portion of the Freeze map whose keys are greater than (or equal to, if `inclusive` is true) the given key. If `inclusive` is not specified, the method behaves as if `inclusive` is true.

- `public NavigableMap<`*`MEMBER_TYPE`*`, Set<Map.Entry<K, V>>>`
  `subMapFor`*`MEMBER`*`(`*`MEMBER_TYPE`*` from, boolean fromInclusive,`
  `MEMBER_TYPE`` to, boolean toInclusive)`

Copyright 2016, ZeroC, Inc.

`public NavigableMap<`*`MEMBER_TYPE`*`, Set<Map.Entry<K, V>>>`
`subMapFor`*`MEMBER`*`(`*`MEMBER_TYPE`*` from, `*`MEMBER_TYPE`*` to)`
Returns a view of the portion of the Freeze map whose keys are within the given range. If `fromInclusive` and `toInclusive` are not specified, the method behaves as if `fromInclusive` is true and `toInclusive` is false.

- `public NavigableMap<`*`MEMBER_TYPE`*`, Set<Map.Entry<K, V>>>`
`mapFor`*`MEMBER`*`()`
Returns a view of the entire Freeze map ordered by the index key.

For the methods returning a `NavigableMap`, the key type is the secondary key type and the value is the set of matching key-value pairs from the Freeze map. (For the sake of readability, we have omitted the `java.util` prefix from `Set` and `Map.Entry`.) In other words, the returned map is a mapping of the secondary key to all of the entries whose value contains the same key. Any attempt to add, remove, or modify an element via a sub map view or an iterator of a sub map view results in an `UnsupportedOperationException`.

Note that iterators returned by the `findBy`*`MEMBER`* methods, as well as those created for sub map views, may need to be closed explicitly, just like iterators obtained for the main Freeze map.

Here are the definitions of the index methods for `StringIntMap`:

Java

```
public Freeze.Map.EntryIterator<Map.Entry<String, Integer>>
findByValue(Integer index);

public Freeze.Map.EntryIterator<Map.Entry<String, Integer>>
findByValue(Integer index, boolean onlyDups);

public int valueCount(Integer index);

public NavigableMap<Integer, Set<Map.Entry<String, Integer>>>
headMapForValue(Integer to, boolean inclusive);
public NavigableMap<Integer, Set<Map.Entry<String, Integer>>>
headMapForValue(Integer to);

public NavigableMap<Integer, Set<Map.Entry<String, Integer>>>
tailMapForValue(Integer from, boolean inclusive);
public NavigableMap<Integer, Set<Map.Entry<String, Integer>>>
tailMapForValue(Integer from);

public NavigableMap<Integer, Set<Map.Entry<String, Integer>>>
subMapForValue(Integer from, boolean fromInclusive,
               Integer to, boolean toInclusive);
public NavigableMap<Integer, Set<Map.Entry<String, Integer>>>
subMapForValue(Integer from, Integer to);

public NavigableMap<Integer, Set<Map.Entry<String, Integer>>>
mapForValue();
```

Sample Freeze Map Program in Java

The program below demonstrates how to use a `StringIntMap` to store `<string, int>` pairs in a database. You will notice that there are no explicit `read` or `write` operations called by the program; instead, simply using the map has the side effect of accessing the database.

Java

Copyright 2016, ZeroC, Inc.

```
public class Client
{
    public static void
    main(String[] args)
    {
        // Initialize the Communicator.
        //
        Ice.Communicator communicator = Ice.Util.initialize(args);

        // Create a Freeze database connection.
        //
        Freeze.Connection connection =
Freeze.Util.createConnection(communicator, "db");

        // Instantiate the map.
        //
        StringIntMap map = new StringIntMap(connection, "simple", true);

        // Clear the map.
        //
        map.clear();

        int i;

        // Populate the map.
        //
        for (i = 0; i < 26; i++) {
            final char[] ch = { (char)('a' + i) };
            map.put(new String(ch), i);
        }

        // Iterate over the map and change the values.
        //
        for (java.util.Map.Entry<String, Integer> e : map.entrySet()) {
            Integer in = e.getValue();
            e.setValue(in.intValue() + 1);
        }

        // Find and erase the last element.
        //
        boolean b;
        b = map.containsKey("z");
        assert(b);
        b = map.fastRemove("z");
        assert(b);

        // Clean up.
        //
        map.close();
        connection.close();
        communicator.destroy();

        System.exit(0);
    }
}
```


Copyright 2016, ZeroC, Inc.

Prior to instantiating a Freeze map, the application must connect to a Berkeley DB database environment:

Java

```
Freeze.Connection connection = Freeze.Util.createConnection(communicator, "db");
```

The second argument is the name of a Berkeley DB database environment; by default, this is also the file system directory in which Berkeley DB creates all database and administrative files.

Next, the code instantiates the `StringIntMap` on the connection. The constructor's second argument supplies the name of the database file, and the third argument indicates that the database should be created if it does not exist:

Java

```
StringIntMap map = new StringIntMap(connection, "simple", true);
```

After instantiating the map, we clear it to make sure it is empty in case the program is run more than once:

Java

```
map.clear();
```

We populate the map, using a single-character string as the key. As with `java.util.Map`, the key and value types must be Java objects but the compiler takes care of autoboxing the integer argument:

Java

```
for (i = 0; i < 26; i++) {
    final char[] ch = { (char)('a' + i) };
    map.put(new String(ch), i);
}
```

Iterating over the map is no different from iterating over any other map that implements the `java.util.Map` interface:

Java

```
for (java.util.Map.Entry<String, Integer> e : map.entrySet()) {
    Integer in = e.getValue();
    e.setValue(in.intValue() + 1);
}
```

Next, the program verifies that an element exists with key `z`, and then removes it using `fastRemove`:

Copyright 2016, ZeroC, Inc.

Java

```
b = map.containsKey("z");
assert(b);
b = map.fastRemove("z");
assert(b);
```

Finally, the program closes the map and its connection.

Java

```
map.close();
connection.close();
```

See Also

- Using the Slice Compilers
- slice2freezej Command-Line Options
- Freeze Map Concepts

Copyright 2016, ZeroC, Inc.

slice2freezej Command-Line Options

The Slice-to-Freeze compiler, `slice2freezej`, creates Java classes for Freeze maps. The compiler offers the following command-line options in addition to the standard options:

--dict NAME,KEY,VALUE

Generate a Freeze map class named *NAME* using *KEY* as key and *VALUE* as value. This option may be specified multiple times to generate several Freeze maps. *NAME* may be a scoped Java name, such as `Demo.Struct1ObjectMap`. *KEY* and *VALUE* represent Slice types and therefore must use Slice syntax, such as `bool` or `Ice::Identity`. The type identified by *KEY* must be a legal dictionary key type.

--dict-index MAP[,MEMBER][,case-sensitive|case-insensitive]

Add an index to the Freeze map named *MAP*. If *MEMBER* is specified, the map value type must be a structure or a class, and *MEMBER* must be the name of a member of that type. If *MEMBER* is not specified, the entire value is indexed. When the indexed member (or entire value) is a string, the index can be case-sensitive (default) or case-insensitive.

--index CLASS,TYPE,MEMBER[,case-sensitive|case-insensitive]

Generate an index class for a Freeze evictor. *CLASS* is the name of the index class to be generated. *TYPE* denotes the type of class to be indexed (objects of different classes are not included in this index). *MEMBER* is the name of the data member in *TYPE* to index. When *MEMBER* has type `string`, it is possible to specify whether the index is case-sensitive or not. The default is case-sensitive.

--meta META

Define the global metadata directive *META*. Using this option is equivalent to defining the global metadata *META* in each named Slice file, as well as in any file included by a named Slice file.

See Also

- Using the Slice Compilers
- slice2java Command-Line Options
- Freeze Map Concepts

Copyright 2016, ZeroC, Inc.

Using a Freeze Map in the File System Server

We can use a Freeze map to add persistence to the file system server, and we'll present implementations in both C++ and Java. However, a Freeze evictor is often a better choice for applications (such as the file system server) in which the persistent value is an Ice object.

In general, incorporating a Freeze map into your application requires the following steps:

1. Evaluate your existing Slice definitions for suitable key and value types.
2. If no suitable key or value types are found, define new (possibly derived) types that capture your persistent state requirements. Consider placing these definitions in a separate file: these types are only used by the server for persistence, and therefore do not need to appear in the "public" definitions required by clients. Also consider placing your persistent types in a separate module to avoid name clashes.
3. Generate a Freeze map for your persistent types using the Freeze compiler.
4. Use the Freeze map in your operation implementations.

Choosing Key and Value Types for the File System

Our goal is to implement the file system using Freeze maps for all persistent storage, including files and their contents. There are various options for how to implement the server. For this example, the server is stateless; whenever a client invokes an operation, the server accesses the database to satisfy the request. Implementing the server in this way has the advantage that it scales very well: we do not need a separate servant for each node; instead two default servants, one for directories and one for files, are sufficient. This keeps the memory requirements of the server to a minimum and also allows us to rely on the database for transactions and locking. (This is a very common implementation technique for servers that act as a front end to a database: the server is a simple facade that implements each operation by accessing the database.)

Our first step is to select the Slice types we will use for the key and value types for our maps. For each file, we need to store the name of the file, its parent directory, and the contents of the file. For directories, we also store the name and parent directory, as well as a dictionary that keeps track of the subdirectories and files in that directory. This leads to Slice definitions (in file `FilesystemDB.ice`) as follows:

Slice

```
#include <Filesystem.ice>
#include <Ice/Identity.ice>

module FilesystemDB {
    struct FileEntry {
        string name;
        Ice::Identity parent;
        Filesystem::Lines text;
    };

    dictionary<string, Filesystem::NodeDesc> StringNodeDescDict;

    struct DirectoryEntry {
        string name;
        Ice::Identity parent;
        StringNodeDescDict nodes;
    };
};
```

Note that the definitions are placed into a separate module, so they do not affect the existing definitions of the non-persistent version of the application. For reference, here is the definition of `NodeDesc` once more:

Copyright 2016, ZeroC, Inc.

Slice

```
module Filesystem {
    // ...

    enum NodeType { DirType, FileType };

    struct NodeDesc {
        string name;
        NodeType type;
        Node* proxy;
    };

    // ...
};
```

To store the persistent state for the file system, we use two Freeze maps: one map for files and one map for directories. For files, we map the identity of the file to its corresponding `FileEntry` structure and, similarly, for directories, we map the identity of the directory to its corresponding `DirectoryEntry` structure.

When a request arrives from a client, the object identity is available in the server. The server uses the identity to retrieve the state of the target node for the request from the database and act on that state accordingly.

Topics

- Adding a Freeze Map to the C++ File System Server
- Adding a Freeze Map to the Java File System Server

See Also

- Default Servants
- Freeze Evictors

Copyright 2016, ZeroC, Inc.

Adding a Freeze Map to the C++ File System Server

Here we present a C++ implementation of the file system server.

On this page:

- Generating the File System Maps in C++
- The Server Main Program in C++
- The Servant Class Definitions in C++
- Implementing FileI with a Freeze Map in C++
- Implementing DirectoryI with a Freeze Map in C++

Generating the File System Maps in C++

Now that we have selected our key and value types, we can generate the maps as follows:

```
$ slice2freeze -I$(ICE_HOME)/slice -I. --ice --dict \
    FilesystemDB::IdentityFileEntryMap,Ice::Identity,\
    FilesystemDB::FileEntry \
    IdentityFileEntryMap FilesystemDB.ice \
    $(ICE_HOME)/slice/Ice/Identity.ice
$ slice2freeze -I$(ICE_HOME)/slice -I. --ice --dict \
    FilesystemDB::IdentityDirectoryEntryMap,Ice::Identity,\
    FilesystemDB::DirectoryEntry \
    IdentityDirectoryEntryMap FilesystemDB.ice \
    $(ICE_HOME)/slice/Ice/Identity.ice
```

The resulting map classes are named `IdentityFileEntryMap` and `IdentityDirectoryEntryMap`.

The Server Main Program in C++

The server's `main` program is very simple:

C++

```
#include <FilesystemI.h>
#include <IdentityFileEntryMap.h>
#include <IdentityDirectoryEntryMap.h>
#include <Ice/Application.h>
#include <Freeze/Freeze.h>

using namespace std;
using namespace Filesystem;
using namespace FilesystemDB;

class FilesystemApp : public virtual Ice::Application
{
public:

    FilesystemApp(const string& envName)
        : _envName(envName)
    {
    }

    virtual int run(int, char*[])
```


Copyright 2016, ZeroC, Inc.

```
    {
        shutdownOnInterrupt();

        Ice::ObjectAdapterPtr adapter =
            communicator()->createObjectAdapter("MapFilesystem");

        const Freeze::ConnectionPtr connection(
            Freeze::createConnection(communicator(), _envName));

        const IdentityFileEntryMap fileDB(connection, FileI::filesDB());
        const IdentityDirectoryEntryMap dirDB(
                                    connection,
                                    DirectoryI::directoriesDB());

        adapter->addDefaultServant(new FileI(communicator(), _envName), "file");
        adapter->addDefaultServant(new DirectoryI(communicator(), _envName), "");

        adapter->activate();

        communicator()->waitForShutdown();

        if(interrupted())
            cerr << appName() << ": received signal, shutting down" << endl;

        return 0;
    }

private:

    string _envName;
};

int
main(int argc, char* argv[])
{
    FilesystemApp app("db");
```


Copyright 2016, ZeroC, Inc.

```
    return app.main(argc, argv, "config.server");
}
```

Let us examine the code in detail. First, we are now including `IdentityFileEntry.h` and `IdentityDirectoryEntry.h`. These header files includes all of the other Freeze (and Ice) header files we need.

Next, we define the class `FilesystemApp` as a subclass of `Ice::Application`, and provide a constructor taking a string argument:

C++

```
FilesystemApp(const string& envName)
    : _envName(envName) {}
```

The string argument represents the name of the database environment, and is saved for later use in `run`.

The interesting part of `run` are the few lines of code that create the database connection and the two maps that store files and directories, plus the code to add the two default servants:

C++

```
    const Freeze::ConnectionPtr connection(
        Freeze::createConnection(communicator(), _envName));

    const IdentityFileEntryMap fileDB(connection, FileI::filesDB());
    const IdentityDirectoryEntryMap dirDB(
                              connection,
                              DirectoryI::directoriesDB());

    adapter->addDefaultServant(new FileI(communicator(), _envName), "file");
    adapter->addDefaultServant(new DirectoryI(communicator(), _envName), "");
```

`run` keeps the database connection open for the duration of the program for performance reasons. As we will see shortly, individual operation implementations will use their own connections; however, it is substantially cheaper to create second (and subsequent connections) than it is to create the first connection.

For the default servants, we use `file` as the category for files. For directories, we use the empty default category.

The Servant Class Definitions in C++

The class definition for `FileI` is very simple:

Copyright 2016, ZeroC, Inc.

C++

```
namespace Filesystem {
    class FileI : public File {
    public:
        FileI(const Ice::CommunicatorPtr& communicator,
              const std::string& envName);

        // Slice operations...

        static std::string filesDB();

    private:
        void halt(const Freeze::DatabaseException& ex) const;

        const Ice::CommunicatorPtr _communicator;
        const std::string _envName;
    };
}
```

The `FileI` class stores the communicator and the environment name. These members are initialized by the constructor. The `filesDB` stati c member function returns the name of the file map, and the `halt` member function is used to stop the server if it encounters a catastrophic error.

The `DirectoryI` class looks very much the same, also storing the communicator and environment name. The `directoriesDB` static member function returns the name of the directory map.

C++

```
namespace Filesystem {
    class DirectoryI : public Directory {
    public:
        DirectoryI(const Ice::CommunicatorPtr& communicator,
                   const std::string& envName);

        // Slice operations...

        static std::string directoriesDB();

    private:
        void halt(const Freeze::DatabaseException& ex) const;

        const Ice::CommunicatorPtr _communicator;
        const std::string _envName;
    };
}
```

Implementing `FileI` with a Freeze Map in C++

The `FileI` constructor and the `filesDB` and `halt` member functions have trivial implementations:

Copyright 2016, ZeroC, Inc.

C++

```
FileI::FileI(const Ice::CommunicatorPtr& communicator,
             const string& envName)
    : _communicator(communicator), _envName(envName)
{
}

string
FileI::filesDB()
{
    return "files";
}

void
FileI::halt(const Freeze::DatabaseException& ex) const
{
    Ice::Error error(_communicator->getLogger());
    error << "fatal exception: " << ex << "\n*** Aborting application ***";

    abort();
}
```

The Slice operations all follow the same implementation strategy: we create a database connection and the file map and place the body of the operation into an infinite loop:

C++

```
string
FileI::someOperation(/* ... */ const Ice::Current& c)
{
    const Freeze::ConnectionPtr connection(
        Freeze::createConnection(_communicator, _envName));
    IdentityFileEntryMap fileDB(connection, filesDB());

    for (;;) {
        try {

            // Operation implementation here...

        } catch (const Freeze::DeadlockException&) {
            continue;
        } catch (const Freeze::DatabaseException& ex) {
            halt(ex);
        }
    }
}
```

Each operation creates its own database connection and map for concurrency reasons: the database takes care of all the necessary locking, so there is no need for any other synchronization in the server. If the database detects a deadlock, the code handles the corresponding `Dea`

Copyright 2016, ZeroC, Inc.

`dlockException` and simply tries again until the operation eventually succeeds; any other database exception indicates that something has gone seriously wrong and terminates the server.

Here is the implementation of the `name` method:

C++

```
string
FileI::name(const Ice::Current& c)
{
    const Freeze::ConnectionPtr connection(
        Freeze::createConnection(_communicator, _envName));
    IdentityFileEntryMap fileDB(connection, filesDB());

    for (;;) {
        try {
            IdentityFileEntryMap::iterator p = fileDB.find(c.id);
            if (p == fileDB.end()) {
                throw Ice::ObjectNotExistException(__FILE__, __LINE__);
            }
            return p->second.name;
        } catch (const Freeze::DeadlockException&) {
            continue;
        } catch (const Freeze::DatabaseException& ex) {
            halt(ex);
        }
    }
}
```

The implementation could hardly be simpler: the default servant uses the identity in the `Current` object to index into the file map. If a record with this identity exists, it returns the name of the file as stored in the `FileEntry` structure in the map. Otherwise, if no such entry exists, it throws `ObjectNotExistException`. This happens if the file existed at some time in the past but has since been destroyed.

The `read` implementation is almost identical. It returns the text that is stored by the `FileEntry`:

Copyright 2016, ZeroC, Inc.

C++

```
Lines
FileI::read(const Ice::Current& c)
{
    const Freeze::ConnectionPtr connection(
        Freeze::createConnection(_communicator, _envName));
    IdentityFileEntryMap fileDB(connection, filesDB());

    for (;;) {
        try {
            IdentityFileEntryMap::iterator p = fileDB.find(c.id);
            if (p == fileDB.end()) {
                throw Ice::ObjectNotExistException(__FILE__, __LINE__);
            }
            return p->second.text;
        } catch (const Freeze::DeadlockException&) {
            continue;
        } catch (const Freeze::DatabaseException& ex) {
            halt(ex);
        }
    }
}
```

The `write` implementation updates the file contents and calls `set` on the iterator to update the map with the new contents:

Copyright 2016, ZeroC, Inc.

C++

```
void
FileI::write(const Filesystem::Lines& text, const Ice::Current& c)
{
    const Freeze::ConnectionPtr connection(
        Freeze::createConnection(_communicator, _envName));
    IdentityFileEntryMap fileDB(connection, filesDB());

    for (;;) {
        try {
            IdentityFileEntryMap::iterator p = fileDB.find(c.id);
            if (p == fileDB.end()) {
                throw Ice::ObjectNotExistException(__FILE__, __LINE__);
            }
            FileEntry entry = p->second;
            entry.text = text;
            p.set(entry);
            break;
        } catch (const Freeze::DeadlockException&) {
            continue;
        } catch (const Freeze::DatabaseException& ex) {
            halt(ex);
        }
    }
}
```

Finally, the `destroy` implementation for files must update two maps: it needs to remove its own entry in the file map as well as update the `nodes` map in the parent to remove itself from the parent's map of children. This raises a potential problem: if one update succeeds but the other one fails, we end up with an inconsistent file system: either the parent still has an entry to a non-existent file, or the parent lacks an entry to a file that still exists.

To make sure that the two updates happen atomically, `destroy` performs them in a transaction:

Copyright 2016, ZeroC, Inc.

C++

```
void
FileI::destroy(const Ice::Current& c)
{
    const Freeze::ConnectionPtr connection(
        Freeze::createConnection(_communicator, _envName));
    IdentityFileEntryMap fileDB(connection, filesDB());
    IdentityDirectoryEntryMap dirDB(connection, DirectoryI::directoriesDB());

    for (;;) {
        try {
            Freeze::TransactionHolder txn(connection);

            IdentityFileEntryMap::iterator p = fileDB.find(c.id);
            if (p == fileDB.end()) {
                throw Ice::ObjectNotExistException(__FILE__, __LINE__);
            }
            FileEntry entry = p->second;

            IdentityDirectoryEntryMap::iterator pp = dirDB.find(entry.parent);
            if (pp == dirDB.end()) {
                halt(Freeze::DatabaseException(
                       __FILE__, __LINE__,
                       "consistency error: file without parent"));
            }

            DirectoryEntry dirEntry = pp->second;
            dirEntry.nodes.erase(entry.name);
            pp.set(dirEntry);

            fileDB.erase(p);
            txn.commit();
            break;
        } catch (const Freeze::DeadlockException&) {
            continue;
        } catch (const Freeze::DatabaseException& ex) {
            halt(ex);
        }
    }
}
```

As you can see, the code first establishes a transaction and then locates the file in the parent directory's map of nodes. After removing the file from the parent, the code updates the parent's persistent state by calling `set` on the parent iterator and then removes the file from the file map before committing the transaction.

Implementing `DirectoryI` with a Freeze Map in C++

The `DirectoryI::directoriesDB` implementation returns the string `directories`, and the `halt` implementation is the same as for `FileI`, so we do not show them here.

Turning to the constructor, we must cater for two different scenarios:

- The server is started with a database that already contains a number of nodes.

Copyright 2016, ZeroC, Inc.

- The server is started for the very first time with an empty database.

This means that the root directory (which must always exist) may or may not be present in the database. Accordingly, the constructor looks for the root directory (with the fixed identity `RootDir`); if the root directory does not exist in the database, it creates it:

C++

```
DirectoryI::DirectoryI(const Ice::CommunicatorPtr& communicator,
const string& envName)
    : _communicator(communicator), _envName(envName)
{
    const Freeze::ConnectionPtr connection =
        Freeze::createConnection(_communicator, _envName);
    IdentityDirectoryEntryMap dirDB(connection, directoriesDB());

    for (;;) {
        try {
            Ice::Identity rootId;
            rootId.name = "RootDir";
            IdentityDirectoryEntryMap::const_iterator p = dirDB.find(rootId);
            if (p == dirDB.end()) {
                DirectoryEntry d;
                d.name = "/";
                dirDB.put(make_pair(rootId, d));
            }
            break;
        } catch (const Freeze::DeadlockException&) {
            continue;
        } catch (const Freeze::DatabaseException& ex) {
            halt(ex);
        }
    }
}
```

Next, let us examine the implementation of `createDirectory`. Similar to the `FileI::destroy` operation, `createDirectory` must update both the parent's nodes map and create a new entry in the directory map. These updates must happen atomically, so we perform them in a separate transaction:

Copyright 2016, ZeroC, Inc.

C++

```
DirectoryPrx
DirectoryI::createDirectory(const string& name, const Ice::Current& c)
{
    const Freeze::ConnectionPtr connection(
        Freeze::createConnection(_communicator, _envName));
    IdentityDirectoryEntryMap directoryDB(connection, directoriesDB());

    for (;;) {
        try {
            Freeze::TransactionHolder txn(connection);

            IdentityDirectoryEntryMap::iterator p =
                directoryDB.find(c.id);
            if (p == directoryDB.end()) {
                throw Ice::ObjectNotExistException(__FILE__, __LINE__);
            }

            DirectoryEntry entry = p->second;
            if (name.empty() || entry.nodes.find(name) != entry.nodes.end()) {
                throw NameInUse(name);
            }

            DirectoryEntry d;
            d.name = name;
            d.parent = c.id;

            Ice::Identity id;
            id.name = IceUtil::generateUUID();
            DirectoryPrx proxy = DirectoryPrx::uncheckedCast(c.adapter->createProxy
(id));

            NodeDesc nd;
            nd.name = name;
            nd.type = DirType;
            nd.proxy = proxy;
            entry.nodes.insert(make_pair(name, nd));

            p.set(entry);
            directoryDB.put(make_pair(id, d));

            txn.commit();

            return proxy;
        } catch (const Freeze::DeadlockException&) {
            continue;
        } catch (const Freeze::DatabaseException& ex) {
            halt(ex);
        }
    }
}
```


Copyright 2016, ZeroC, Inc.

After establishing the transaction, the code ensures that the directory does not already contain an entry with the same name and then initializes a new `DirectoryEntry`, setting the name to the name of the new directory, and the parent to its own identity. The identity of the new directory is a UUID, which ensures that all directories have unique identities. In addition, the UUID prevents the accidental rebirth of a file or directory in the future.

The code then initializes a new `NodeDesc` structure with the details of the new directory and, finally, updates its own map of children as well as adding the new directory to the map of directories before committing the transaction.

The `createFile` implementation is almost identical, so we do not show it here. Similarly, the `name` and `destroy` implementations are almost identical to the ones for `FileI`, so let us move to `list`:

C++

```
NodeDescSeq
DirectoryI::list(const Ice::Current& c)
{
    const Freeze::ConnectionPtr connection(
        Freeze::createConnection(_communicator, _envName));
    IdentityDirectoryEntryMap directoryDB(connection, directoriesDB());

    for (;;) {
        try {
            IdentityDirectoryEntryMap::iterator p = directoryDB.find(c.id);
            if (p == directoryDB.end()) {
                throw Ice::ObjectNotExistException(__FILE__, __LINE__);
            }
            NodeDescSeq result;
            for (StringNodeDescDict::const_iterator q = p->second.nodes.begin();
                 q != p->second.nodes.end(); ++q) {
                result.push_back(q->second);
            }
            return result;
        } catch (const Freeze::DeadlockException&) {
            continue;
        } catch (const Freeze::DatabaseException& ex) {
            halt(ex);
        }
    }
}
```

Again, the code is very simple: it iterates over the `nodes` map, adding each `NodeDesc` structure to the returned sequence.

The `find` implementation is even simpler, so we do not show it here.

See Also

- Freeze Maps
- Object Identity and Uniqueness
- The Current Object

Copyright 2016, ZeroC, Inc.

Adding a Freeze Map to the Java File System Server

Here we present a Java implementation of the file system server.

On this page:

- Generating the File System Maps in Java
- The Server Main Program in Java
- Implementing FileI with a Freeze Map in Java
- Implementing DirectoryI with a Freeze Map in Java

Generating the File System Maps in Java

Now that we have selected our key and value types, we can generate the maps as follows:

```
$ slice2freezej -I$(ICE_HOME)/slice -I. --ice --dict \
    FilesystemDB.IdentityFileEntryMap,Ice::Identity,\
    FilesystemDB::FileEntry \
    IdentityFileEntryMap FilesystemDB.ice \
    $(ICE_HOME)/slice/Ice/Identity.ice
$ slice2freezej -I$(ICE_HOME)/slice -I. --ice --dict \
    FilesystemDB.IdentityDirectoryEntryMap,Ice::Identity,\
    FilesystemDB::DirectoryEntry \
    IdentityDirectoryEntryMap FilesystemDB.ice \
    $(ICE_HOME)/slice/Ice/Identity.ice
```

The resulting map classes are named `IdentityFileEntryMap` and `IdentityDirectoryEntryMap`.

The Server Main Program in Java

The server's main program is very simple:

Copyright 2016, ZeroC, Inc.

Java

```
import Filesystem.*;
import FilesystemDB.*;

public class Server extends Ice.Application
{
    public
    Server(String envName)
    {
        _envName = envName;
    }

    public int
    run(String[] args)
    {
        Ice.ObjectAdapter adapter =
communicator().createObjectAdapter("MapFilesystem");

        Freeze.Connection connection = null;
        try {
            connection = Freeze.Util.createConnection(communicator(), _envName);
            IdentityFileEntryMap fileDB =
                new IdentityFileEntryMap(connection, FileI.filesDB(), true);
            IdentityDirectoryEntryMap dirDB =
                new IdentityDirectoryEntryMap(connection, DirectoryI.directoriesDB(
), true);

            adapter.addDefaultServant(new FileI(communicator(), _envName), "file");
            adapter.addDefaultServant(new DirectoryI(communicator(), _envName), "")
;


            adapter.activate();

            communicator().waitForShutdown();
        } finally {
            connection.close();
        }

        return 0;
    }

    public static void
    main(String[] args)
    {
        Server app = new Server("db");
        app.main("MapServer", args, "config.server");
        System.exit(0);
    }

    private String _envName;
}
```


Copyright 2016, ZeroC, Inc.

First, we import the `Filesystem` and `FilesystemDB` packages.

Next, we define the class `FilesystemApp` as a subclass of `Ice.Application`, and provide a constructor taking a string argument:

Java

```
    FilesystemApp(const string& envName)
        : _envName(envName) {}
```

The string argument represents the name of the database environment, and is saved for later use in `run`.

The interesting part of `run` are the few lines of code that create the database connection and the two maps that store files and directories, plus the code to add the two default servants:

Java

```
            connection = Freeze.Util.createConnection(communicator(), _envName);
            IdentityFileEntryMap fileDB =
                new IdentityFileEntryMap(connection, FileI.filesDB(), true);
            IdentityDirectoryEntryMap dirDB =
                new IdentityDirectoryEntryMap(connection, DirectoryI.directoriesDB(
), true);

            adapter.addDefaultServant(new FileI(communicator(), _envName), "file");
            adapter.addDefaultServant(new DirectoryI(communicator(), _envName), "")
;
```

`run` keeps the database connection open for the duration of the program for performance reasons. As we will see shortly, individual operation implementations will use their own connections; however, it is substantially cheaper to create second (and subsequent connections) than it is to create the first connection.

For the default servants, we use `file` as the category for files. For directories, we use the empty default category.

Implementing `FileI` with a Freeze Map in Java

The class definition for `FileI` is very simple:

Copyright 2016, ZeroC, Inc.

Java

```
public class FileI extends _FileDisp
{
    public
    FileI(Ice.Communicator communicator, String envName)
    {
        _communicator = communicator;
        _envName = envName;
    }

    // Slice operations...

    public static String
    filesDB()
    {
        return "files";
    }

    private void
    halt(Freeze.DatabaseException e)
    {
        java.io.StringWriter sw = new java.io.StringWriter();
        java.io.PrintWriter pw = new java.io.PrintWriter(sw);
        e.printStackTrace(pw);
        pw.flush();
        _communicator.getLogger().error(
            "fatal database error\n" + sw.toString() + "\n*** Halting JVM ***");
        Runtime.getRuntime().halt(1);
    }

    private Ice.Communicator _communicator;
    private String _envName;
}
```

The `FileI` class stores the communicator and the environment name. These members are initialized by the constructor. The `filesDB` static method returns the name of the file map, and the `halt` member function is used to stop the server if it encounters a catastrophic error.

The Slice operations all follow the same implementation strategy: we create a database connection and the file map and place the body of the operation into an infinite loop:

Copyright 2016, ZeroC, Inc.

Java

```
    public String
    someOperation(/* ... */ Ice.Current c)
    {
        Freeze.Connection connection =
Freeze.Util.createConnection(_communicator, _envName);
        try {
            IdentityFileEntryMap fileDB =
new IdentityFileEntryMap(connection, filesDB());

            for (;;) {
                try {

                    // Operation implementation here...

                } catch (Freeze.DeadlockException ex) {
                    continue;
                } catch (Freeze.DatabaseException ex) {
                    halt(ex);
                }
            }
        } finally {
            connection.close();
        }
    }
```

Each operation creates its own database connection and map for concurrency reasons: the database takes care of all the necessary locking, so there is no need for any other synchronization in the server. If the database detects a deadlock, the code handles the corresponding `DeadlockException` and simply tries again until the operation eventually succeeds; any other database exception indicates that something has gone seriously wrong and terminates the server.

Here is the implementation of the `name` method:

Copyright 2016, ZeroC, Inc.

Java

```
    public String
    name(Ice.Current c)
    {
        Freeze.Connection connection =
Freeze.Util.createConnection(_communicator, _envName);
        try
        {
            IdentityFileEntryMap fileDB =
new IdentityFileEntryMap(connection, filesDB());

            for (;;) {
                try {
                    FileEntry entry = fileDB.get(c.id);
                    if (entry == null) {
                        throw new Ice.ObjectNotExistException();
                    }
                    return entry.name;
                } catch (Freeze.DeadlockException ex) {
                    continue;
                } catch (Freeze.DatabaseException ex) {
                    halt(ex);
                }
            }
        } finally {
            connection.close();
        }
    }
```

The implementation could hardly be simpler: the default servant uses the identity in the `Current` object to index into the file map. If a record with this identity exists, it returns the name of the file as stored in the `FileEntry` structure in the map. Otherwise, if no such entry exists, it throws `ObjectNotExistException`. This happens if the file existed at some time in the past but has since been destroyed.

The `read` implementation is almost identical. It returns the text that is stored by the `FileEntry`:

Copyright 2016, ZeroC, Inc.

Java

```
    public String[]
    read(Ice.Current c)
    {
        Freeze.Connection connection =
            Freeze.Util.createConnection(_communicator, _envName);
        try {
            IdentityFileEntryMap fileDB = new
IdentityFileEntryMap(connection, filesDB());

            for (;;) {
                try {
                    FileEntry entry = fileDB.get(c.id);
                    if (entry == null) {
                        throw new Ice.ObjectNotExistException();
                    }
                    return entry.text;
                } catch (Freeze.DeadlockException ex) {
                    continue;
                } catch (Freeze.DatabaseException ex) {
                    halt(ex);
                }
            }
        } finally {
            connection.close();
        }
    }
```

The `write` implementation updates the file contents and calls `put` on the iterator to update the map with the new contents:

Copyright 2016, ZeroC, Inc.

Java

```
    public void
    write(String[] text, Ice.Current c)
        throws GenericError
    {
        Freeze.Connection connection =
Freeze.Util.createConnection(_communicator, _envName);
        try {
            IdentityFileEntryMap fileDB =
new IdentityFileEntryMap(connection, filesDB());

            for (;;) {
                try {
                    FileEntry entry = fileDB.get(c.id);
                    if (entry == null) {
                        throw new Ice.ObjectNotExistException();
                    }
                    entry.text = text;
                    fileDB.put(c.id, entry);
                    break;
                } catch (Freeze.DeadlockException ex) {
                    continue;
                } catch (Freeze.DatabaseException ex) {
                    halt(ex);
                }
            }
        } finally {
            connection.close();
        }
    }
```

Finally, the `destroy` implementation for files must update two maps: it needs to remove its own entry in the file map as well as update the `nodes` map in the parent to remove itself from the parent's map of children. This raises a potential problem: if one update succeeds but the other one fails, we end up with an inconsistent file system: either the parent still has an entry to a non-existent file, or the parent lacks an entry to a file that still exists.

To make sure that the two updates happen atomically, `destroy` performs them in a transaction:

Java

Copyright 2016, ZeroC, Inc.

```
    public void
    destroy(Ice.Current c)
        throws PermissionDenied
    {
        Freeze.Connection connection =
Freeze.Util.createConnection(_communicator, _envName);
        try {
            IdentityFileEntryMap fileDB =
new IdentityFileEntryMap(connection, filesDB());
            IdentityDirectoryEntryMap dirDB =
                new IdentityDirectoryEntryMap(connection, DirectoryI.directoriesDB(
));

            for (;;) {
                Freeze.Transaction txn = null;
                try {
                    txn = connection.beginTransaction();

                    FileEntry entry = fileDB.get(c.id);
                    if (entry == null) {
                        throw new Ice.ObjectNotExistException();
                    }

                    DirectoryEntry dirEntry = (DirectoryEntry)dirDB.get(entry.paren
t);
                    if (dirEntry == null) {
                        halt(new Freeze.DatabaseException(
                                 "consistency error: file without parent"));
                    }

                    dirEntry.nodes.remove(entry.name);
                    dirDB.put(entry.parent, dirEntry);

                    fileDB.remove(c.id);

                    txn.commit();
                    txn = null;
                    break;
                } catch (Freeze.DeadlockException ex) {
                    continue;
                } catch (Freeze.DatabaseException ex) {
                    halt(ex);
                } finally {
                    if (txn != null) {
                        txn.rollback();
                    }
                }
            }
        }
        finally
        {
            connection.close();
        }
    }
```


Copyright 2016, ZeroC, Inc.

As you can see, the code first establishes a transaction and then locates the file in the parent directory's map of nodes. After removing the file from the parent, the code updates the parent's persistent state by calling `put` on the parent iterator and then removes the file from the file map before committing the transaction.

Implementing `DirectoryI` with a Freeze Map in Java

The `DirectoryI.directoriesDB` implementation returns the string `directories`, and the `halt` implementation is the same as for `FileI`, so we do not show them here.

Turning to the constructor, we must cater for two different scenarios:

- The server is started with a database that already contains a number of nodes.
- The server is started for the very first time with an empty database.

This means that the root directory (which must always exist) may or may not be present in the database. Accordingly, the constructor looks for the root directory (with the fixed identity `RootDir`); if the root directory does not exist in the database, it creates it:

Java

```
    public
    DirectoryI(Ice.Communicator communicator, String envName)
    {
        _communicator = communicator;
        _envName = envName;

        Freeze.Connection connection =
Freeze.Util.createConnection(_communicator, _envName);
        try {
            IdentityDirectoryEntryMap dirDB =
                new IdentityDirectoryEntryMap(connection, directoriesDB());

            for (;;) {
                try {
                    Ice.Identity rootId = new Ice.Identity("RootDir", "");
                    DirectoryEntry entry = dirDB.get(rootId);
                    if (entry == null) {
                        dirDB.put(rootId, new DirectoryEntry("/",
new Ice.Identity("", ""), null));
                    }
                    break;
                } catch (Freeze.DeadlockException ex) {
                    continue;
                } catch (Freeze.DatabaseException ex) {
                    halt(ex);
                }
            }
        } finally {
            connection.close();
        }
    }
```

Next, let us examine the implementation of `createDirectory`. Similar to the `FileI::destroy` operation, `createDirectory` must update both the parent's nodes map and create a new entry in the directory map. These updates must happen atomically, so we perform them in a separate transaction:

Copyright 2016, ZeroC, Inc.

Java

```
    public DirectoryPrx
    createDirectory(String name, Ice.Current c)
        throws NameInUse
    {
        Freeze.Connection connection =
Freeze.Util.createConnection(_communicator, _envName);
        try {
            IdentityDirectoryEntryMap dirDB =
                new IdentityDirectoryEntryMap(connection, directoriesDB());

            for (;;) {
                Freeze.Transaction txn = null;
                try {
                    txn = connection.beginTransaction();

                    DirectoryEntry entry = dirDB.get(c.id);
                    if (entry == null) {
                        throw new Ice.ObjectNotExistException();
                    }
                    if(name.length() == 0 || entry.nodes.get(name) != null) {
                        throw new NameInUse(name);
                    }

                    DirectoryEntry newEntry = new DirectoryEntry(name, c.id, null);
                    Ice.Identity id = new Ice.Identity(java.util.UUID.randomUUID().
toString(), "");
                    DirectoryPrx proxy =
DirectoryPrxHelper.uncheckedCast(c.adapter.createProxy(id));

                    entry.nodes.put(name, new NodeDesc(name, NodeType.DirType,
proxy));
                    dirDB.put(c.id, entry);
                    dirDB.put(id, newEntry);

                    txn.commit();
                    txn = null;

                    return proxy;
                } catch (Freeze.DeadlockException ex) {
                    continue;
                } catch (Freeze.DatabaseException ex) {
                    halt(ex);
                } finally {
                    if(txn != null) {
                        txn.rollback();
                    }
                }
            }
        } finally {
            connection.close();
        }
    }
```


Copyright 2016, ZeroC, Inc.

After establishing the transaction, the code ensures that the directory does not already contain an entry with the same name and then initializes a new `DirectoryEntry`, setting the name to the name of the new directory, and the parent to its own identity. The identity of the new directory is a UUID, which ensures that all directories have unique identities. In addition, the UUID prevents the accidental rebirth of a file or directory in the future.

The code then initializes a new `NodeDesc` structure with the details of the new directory and, finally, updates its own map of children as well as adding the new directory to the map of directories before committing the transaction.

The `createFile` implementation is almost identical, so we do not show it here. Similarly, the `name` and `destroy` implementations are almost identical to the ones for `FileI`, so let us move to `list`:

Java

```
    public NodeDesc[]
    list(Ice.Current c)
    {
        Freeze.Connection connection =
            Freeze.Util.createConnection(_communicator, _envName);
        try {
            IdentityDirectoryEntryMap dirDB =
                new IdentityDirectoryEntryMap(connection, directoriesDB());

            for (;;) {
                try {
                    DirectoryEntry entry = dirDB.get(c.id);
                    if (entry == null) {
                        throw new Ice.ObjectNotExistException();
                    }
                    NodeDesc[] result = new NodeDesc[entry.nodes.size()];
                    java.util.Iterator<NodeDesc> p =
entry.nodes.values().iterator();
                    for (int i = 0; i < entry.nodes.size(); ++i) {
                        result[i] = p.next();
                    }
                    return result;
                } catch (Freeze.DeadlockException ex) {
                    continue;
                } catch (Freeze.DatabaseException ex) {
                    halt(ex);
                }
            }
        } finally {
            connection.close();
        }
    }
```

Again, the code is very simple: it iterates over the `nodes` map, adding each `NodeDesc` structure to the returned sequence.

The `find` implementation is even simpler, so we do not show it here.

See Also

- Freeze Maps
- The Current Object
- Object Identity and Uniqueness

Copyright 2016, ZeroC, Inc.

Freeze Catalogs

In each database environment, Freeze maintains an internal table that contains type information describing all the databases in the environment. This table is an instance of a Freeze map in which the key is a string representing the database name and the value is an instance of `Freeze::CatalogData`:

Slice

```
module Freeze {
    struct CatalogData {
        bool evictor;
        string key;
        string value;
    };
};
```

An entry describes an evictor database if the `evictor` member is true, in which case the `key` and `value` members are empty strings. An entry that describes a Freeze map sets `evictor` to false; the `key` and `value` members contain the Slice types used when the map was defined.

FreezeScript tools such as `transformdb` and `dumpdb` access the catalog to obtain type information when none is supplied by the user. You can also use `dumpdb` to display the catalog of a database environment.

Freeze applications may access the catalog in the same manner as any other Freeze map. For example, the following C++ code displays the contents of a catalog:

C++

```
#include <Freeze/Catalog.h>
...
string envName = ...;
Freeze::ConnectionPtr conn = Freeze::createConnection(communicator, envName);
Freeze::Catalog catalog(conn, Freeze::catalogName());
for (Freeze::Catalog::const_iterator p = catalog.begin();
    p != catalog.end(); ++p) {
    if (p->second.evictor)
        cout << p->first << ": evictor" << endl;
    else
        cout << p->first << ": map<" << p->second.key
             << ", " << p->second.value << ">" << endl;
}
conn->close();
```

The equivalent Java code is shown below:

Copyright 2016, ZeroC, Inc.

Java

```
String envName = ...;
Freeze.Connection conn = Freeze.Util.createConnection(communicator, envName);
Freeze.Catalog catalog = new Freeze.Catalog(conn, Freeze.Util.catalogName(), true);
for (java.util.Map.Entry<String, Freeze.CatalogData> e :
     catalog.entrySet()) {
    String name = e.getKey();
    Freeze.CatalogData data = e.getValue();
    if (data.evictor)
        System.out.println(name + ": evictor");
    else
        System.out.println(name + ": map<" + data.key + ", " + data.value + ">");
}
conn.close();
```

See Also

- Freeze Maps
- Freeze Evictors
- FreezeScript

Copyright 2016, ZeroC, Inc.

Backing Up Freeze Databases

When you store important information in a Freeze database environment, you should consider regularly backing up the database environment.

There are two forms of backups: cold backups, where you just copy your database environment directory while no application is using these files (very straightforward), and hot backups, where you backup a database environment while an application is actively reading and writing data.

In order to perform a hot backup on a Freeze environment, you need to configure this Freeze environment with two non-default settings:

- `Freeze.DbEnv.envName.OldLogsAutoDelete=0`
  This instructs Freeze to keep old log files instead of periodically deleting them. This setting is necessary for proper hot backups; it implies that you will need to take care of deleting old files yourself (typically as part of your periodic backup procedure).
- `Freeze.DbEnv.envName.DbPrivate=0`
  By default, Freeze is configured with `DbPrivate` set to 1, which means only one process at a time can safely access the database environment. When performing hot backups, you need to access this database environment concurrently from various Berkeley DB utilities (such as `db_archive` or `db_hotbackup`), so you need to set this property to 0.

The `Freeze/backup` C++ demo in your Ice distribution shows one way to perform such backups and recovery. Please consult the Berkeley DB documentation for further details.

See Also

- Freeze

Copyright 2016, ZeroC, Inc.

FreezeScript

Freeze supplies a valuable set of services for simplifying the use of persistence in Ice applications. However, while Freeze makes it easy for an application to manage its persistent state, there are additional administrative responsibilities that must also be addressed:

- Migration
  As an application evolves, it is not unusual for the types describing its persistent state to evolve as well. When these changes occur, a great deal of time can be saved if existing databases can be migrated to the new format while preserving as much information as possible.
- Inspection
  The ability to examine a database can be helpful during every stage of the application's lifecycle, from development to deployment.

FreezeScript provides tools for performing both of these activities on Freeze map and evictor databases. These databases have a well-defined structure because the key and value of each record consist of the marshaled bytes of their respective Slice types. This design allows the FreezeScript tools to operate on any Freeze database using only the Slice definitions for the database types.

Topics

- Migrating a Freeze Database
- Inspecting a Freeze Database
- FreezeScript Descriptor Expression Language

Copyright 2016, ZeroC, Inc.

Migrating a Freeze Database

The FreezeScript tool `transformdb` migrates a database created by a Freeze map or evictor. It accomplishes this by comparing the "old" Slice definitions (i.e., the ones that describe the current contents of the database) with the "new" Slice definitions, and making whatever modifications are necessary to ensure that the transformed database is compatible with the new definitions.

This would be difficult to achieve by writing a custom transformation program because that program would require static knowledge of the old and new types, which frequently define many of the same symbols and would therefore prevent the program from being loaded. The `transformdb` tool avoids this issue using an interpretive approach: the Slice definitions are parsed and used to drive the migration of the database records.

The tool supports two modes of operation:

1. Automatic migration – the database is migrated in a single step using only the default set of transformations.
2. Custom migration – you supply a script to augment or override the default transformations.

Topics

- Automatic Database Migration
- Custom Database Migration
- FreezeScript Transformation XML Reference
- Using transformdb

See Also

- Freeze Maps
- Freeze Evictors

Copyright 2016, ZeroC, Inc.

Automatic Database Migration

On this page:

- Type Compatibility Rules for Automatic Migration
- Default Values for Automatic Migration
- Running an Automatic Migration

The default transformations performed by `transformdb` preserve as much information as possible. However, there are practical limits to the tool's capabilities, since the only information it has is obtained by performing a comparison of the Slice definitions.

For example, suppose our old definition for a structure is the following:

Slice

```
struct AStruct {
    int i;
};
```

We want to migrate instances of this struct to the following revised definition:

Slice

```
struct AStruct {
    int j;
};
```

As the developers, we know that the `int` member has been renamed from `i` to `j`, but to `transformdb` it appears that member `i` was removed and member `j` was added. The default transformation results in exactly that behavior: the value of `i` is lost, and `j` is initialized to a default value. If we need to preserve the value of `i` and transfer it to `j`, then we need to use custom migration.

The changes that occur as a type system evolves can be grouped into three categories:

- Data members
  The data members of class and structure types are added, removed, or renamed. As discussed above, the default transformations initialize new and renamed data members to default values.
- Type names
  Types are added, removed, or renamed. New types do not pose a problem for database migration when used to define a new data member; the member is initialized with default values as usual. On the other hand, if the new type replaces the type of an existing data member, then type compatibility becomes a factor (see the following item).

  Removed types generally do not cause problems either, because any uses of that type must have been removed from the new Slice definitions (e.g., by removing data members of that type). There is one case, however, where removed types become an issue, and that is for polymorphic classes.

  Renamed types are a concern, just like renamed data members, because of the potential for losing information during migration. This is another situation for which custom migration is recommended.
- Type content
  Examples of changes of type content include the key type of a dictionary, the element type of a sequence, or the type of a data member. If the old and new types are not compatible, then the default transformation emits a warning, discards the current value, and reinitializes it with a default value.

Type Compatibility Rules for Automatic Migration

Changes in the type of a value are restricted to certain sets of compatible changes. This section describes the type changes supported by the default transformations. All incompatible type changes result in a warning indicating that the current value is being discarded and a default value for the new type assigned in its place. Additional flexibility is provided by custom migration.

Copyright 2016, ZeroC, Inc.

Boolean

A value of type `bool` can be transformed to and from `string`. The legal string values for a `bool` value are `"true"` and `"false"`.

Integer

The integer types `byte`, `short`, `int`, and `long` can be transformed into each other, but only if the current value is within range of the new type. These integer types can also be transformed into `string`.

Floating Point

The floating-point types `float` and `double` can be transformed into each other, as well as to `string`. No attempt is made to detect a loss of precision during transformation.

String

A `string` value can be transformed into any of the primitive types, as well as into enumeration and proxy types, but only if the value is a legal string representation of the new type. For example, the string value `"Pear"` can be transformed into the enumeration `Fruit`, but only if `Pear` is an enumerator of `Fruit`.

Enum

An enumeration can be transformed into an enumeration with the same type ID, or into a string. Transformation between enumerations is performed symbolically. For example, consider our old type below:

Slice

```
enum Fruit { Apple, Orange, Pear };
```

Suppose the enumerator `Pear` is being transformed into the following new type:

Slice

```
enum Fruit { Apple, Pear };
```

The transformed value in the new enumeration is also `Pear`, despite the fact that `Pear` has changed positions in the new type. However, if the old value had been `Orange`, then the default transformation emits a warning because that enumerator no longer exists, and initializes the new value to `Apple` (the default value).

If an enumerator has been renamed, then custom migration is required to convert enumerators from the old name to the new one.

Sequence

A sequence can be transformed into another sequence type, even if the new sequence type does not have the same type id as the old type, but only if the element types are compatible. For example, `sequence<short>` can be transformed into `sequence<int>`, regardless of the names given to the sequence types.

Dictionary

A dictionary can be transformed into another dictionary type, even if the new dictionary type does not have the same type ID as the old type, but only if the key and value types are compatible. For example, `dictionary<int, string>` can be transformed into `dictionary<long, string>`, regardless of the names given to the dictionary types.

Caution is required when changing the key type of a dictionary, because the default transformation of keys could result in duplication. For example, if the key type changes from `int` to `short`, any `int` value outside the range of `short` results in the key being initialized to a

Copyright 2016, ZeroC, Inc.

default value (namely zero). If zero is already used as a key in the dictionary, or another out-of-range key is encountered, then a duplication occurs. The transformation handles key duplication by removing the duplicate element from the transformed dictionary. (Custom migration can be useful in these situations if the default behavior is not acceptable.)

Structure

A `struct` type can only be transformed into another `struct` type with the same type ID. Data members are transformed as appropriate for their types.

Proxy

A proxy value can be transformed into another proxy type, or into `string`. Transformation into another proxy type is done with the same semantics as in a language mapping: if the new type does not match the old type, then the new type must be a base type of the old type (that is, the proxy is widened).

Class

A `class` type can only be transformed into another `class` type with the same type ID. A data member of a `class` type is allowed to be widened to a base type. Data members are transformed as appropriate for their types. See Transforming Objects for more information on transforming classes.

Default Values for Automatic Migration

Data types are initialized with default values, as shown.

Type	Default Value
Boolean	`false`
Numeric	Zero (`0`)
String	Empty string
Enumeration	The first enumerator
Sequence	Empty sequence
Dictionary	Empty dictionary
Struct	Data members are initialized recursively
Proxy	Nil
Class	Nil

Running an Automatic Migration

In order to use automatic transformation, we need to supply the following information to `transformdb`:

- The old and new Slice definitions
- The old and new types for the database key and value
- The database environment directory, the database file name, and the name of a new database environment directory to hold the transformed database

Here is an example of a `transformdb` command:

```
$ transformdb --old old/MyApp.ice --new new/MyApp.ice \
--key int,string --value ::Employee db emp.db newdb
```

Briefly, the `--old` and `--new` options specify the old and new Slice definitions, respectively. These options can be specified as many times as necessary in order to load all of the relevant definitions. The `--key` option indicates that the database key is evolving from `int` to `string`. The `--value` option specifies that `::Employee` is used as the database value type in both old and new type definitions, and therefore only needs to be specified once. Finally, we provide the pathname of the database environment directory (`db`), the file name of the database (`emp.db`), and the pathname of the database environment directory for the transformed database (`newdb`).
See Also

Copyright 2016, ZeroC, Inc.

- Custom Database Migration
- Type IDs
- Using transformdb

Copyright 2016, ZeroC, Inc.

Custom Database Migration

Custom migration is useful when your types have changed in ways that make automatic migration difficult or impossible. It is also convenient to use custom migration when you have complex initialization requirements for new types or new data members, because custom migration enables you to perform many of the same tasks that would otherwise require you to write a throwaway program.

Custom migration operates in conjunction with automatic migration, allowing you to inject your own transformation rules at well-defined intercept points in the automatic migration process. These rules are called *transformation descriptors*, and are written in XML.

On this page:

- Simple Example of Custom Migration
- Overview of Transformation Descriptors
- Transformation Flow of Execution
- Transformation Descriptor Scopes
- Guidelines for Transformation Descriptors

Simple Example of Custom Migration

We can use a simple example to demonstrate the utility of custom migration. Suppose our application uses a Freeze map whose key type is `string` and whose value type is an enumeration, defined as follows:

Slice

```
enum BigThree { Ford, DaimlerChrysler, GeneralMotors };
```

We now wish to rename the enumerator `DaimlerChrysler`, as shown in our new definition:

Slice

```
enum BigThree { Ford, Daimler, GeneralMotors };
```

According to the rules for default transformations, all occurrences of the `DaimlerChrysler` enumerator would be transformed into `Ford`, because `Chrysler` no longer exists in the new definition and therefore the default value `Ford` is used instead.

To remedy this situation, we use the following transformation descriptors:

XML

```
<transformdb>
    <database key="string" value="::BigThree">
        <record>
            <if test="oldvalue == ::Old::DaimlerChrysler">
                <set target="newvalue" value="::New::Daimler"/>
            </if>
        </record>
    </database>
</transformdb>
```

When executed, these descriptors convert occurrences of `DaimlerChrysler` in the old type system into `Daimler` in the transformed database's new type system.

Copyright 2016, ZeroC, Inc.

Overview of Transformation Descriptors

As we saw in the previous example, FreezeScript transformation descriptors are written in XML.

A transformation descriptor file has a well-defined structure. The top-level descriptor in the file is `<transformdb>`. A `<database>` descriptor must be present within `<transformdb>` to define the key and value types used by the database. Inside `<database>`, the `<record>` descriptor triggers the transformation process.

During transformation, type-specific actions are supported by the `<transform>` and `<init>` descriptors, both of which are children of `<transformdb>`. One `<transform>` descriptor and one `<init>` descriptor may be defined for each type in the new Slice definitions. Each time `transformdb` creates a new instance of a type, it executes the `<init>` descriptor for that type, if one is defined. Similarly, each time `transformdb` transforms an instance of an old type into a new type, the `<transform>` descriptor for the new type is executed.

The `<database>`, `<record>`, `<transform>`, and `<init>` descriptors may contain general-purpose action descriptors such as `<if>`, `<set>`, and `<echo>`. These actions resemble statements in programming languages like C++ and Java, in that they are executed in the order of definition and their effects are cumulative. Actions can make use of the expression language that should look familiar to C++ and Java programmers.

Transformation Flow of Execution

The transformation descriptors are executed as follows:

- `<database>` is executed first. Each child descriptor of `<database>` is executed in the order of definition. If a `<record>` descriptor is present, database transformation occurs at that point. Any child descriptors of `<database>` that follow `<record>` are not executed until transformation completes.
- During transformation of each record, `transformdb` creates instances of the new key and value types, which includes the execution of the `<init>` descriptors for those types. Next, the old key and value are transformed into the new key and value, in the following manner:
    1. Locate the `<transform>` descriptor for the type.
    2. If no descriptor is found, or the descriptor exists and it does not preclude default transformation, then transform the data as in automatic database migration.
    3. If the `<transform>` descriptor exists, execute it.
    4. Finally, execute the child descriptors of `<record>`.

Transformation Descriptor Scopes

The `<database>` descriptor creates a global scope, allowing child descriptors of `<database>` to define symbols that are accessible in any descriptor.

In order for a global symbol to be available to a `<transform>` or `<init>` descriptor, the symbol must be defined before the `<record>` descriptor is executed.

Furthermore, certain other descriptors create local scopes that exist only for the duration of the descriptor's execution. For example, the `<transform>` descriptor creates a local scope and defines the symbols `old` and `new` to represent a value in its old and new forms. Child descriptors of `<transform>` can also define new symbols in the local scope, as long as those symbols do not clash with an existing symbol in that scope. It is legal to add a new symbol with the same name as a symbol in an outer scope, but the outer symbol will not be accessible during the descriptor's execution.

The global scope is useful in many situations. For example, suppose you want to track the number of times a certain value was encountered during transformation. This can be accomplished as shown below:

Copyright 2016, ZeroC, Inc.

XML

```
<transformdb>
    <database key="string" value="::Ice::Identity">
        <define name="categoryCount" type="int" value="0"/>
        <record/>
        <echo message="categoryCount = " value="categoryCount"/>
    </database>
    <transform type="::Ice::Identity">
        <if test="new.category == 'Accounting'">
            <set target="categoryCount" value="categoryCount + 1"/>
        </if>
    </transform>
</transformdb>
```

In this example, the `<define>` descriptor introduces the symbol `categoryCount` into the global scope, defining it as type `int` with an initial value of zero. Next, the `<record>` descriptor causes transformation to proceed. Each occurrence of the type `Ice::Identity` causes its `<transform>` descriptor to be executed, which examines the `category` member and increases `categoryCount` if necessary. Finally, after transformation completes, the `<echo>` descriptor displays the final value of `categoryCount`.

To reinforce the relationships between descriptors and scopes, consider the following diagram. Several descriptors are shown, including the symbols they define in their local scopes. In this example, the `<iterate>` descriptor has a dictionary target and therefore the default symbol for the element value, `value`, hides the symbol of the same name in the parent `<init>` descriptor's scope.

This situation can be avoided by assigning a different symbol name to the element value.

In addition to symbols in the `<iterate>` scope, child descriptors of `<iterate>` can also refer to symbols from the `<init>` and `<database>` scopes.

Copyright 2016, ZeroC, Inc.

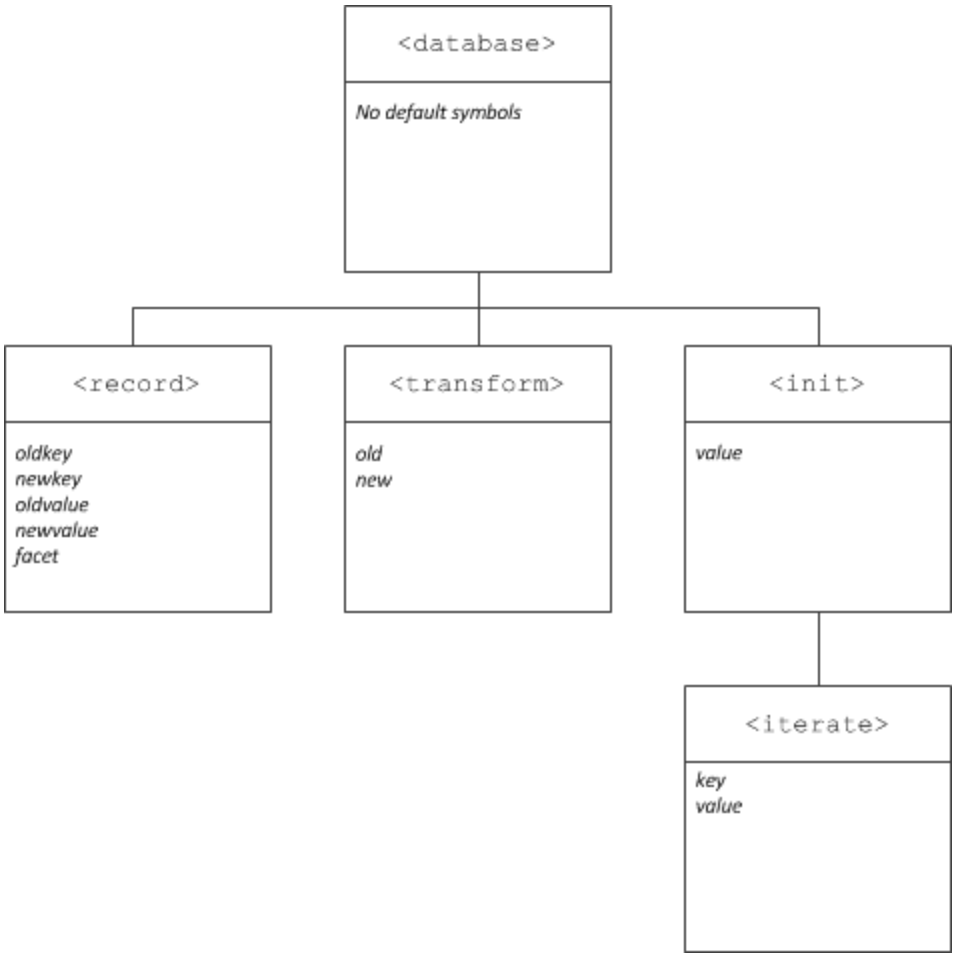

Relationship between descriptors and scopes.

Guidelines for Transformation Descriptors

There are three points at which you can intercept the transformation process: when transforming a record (`<record>`), when transforming an instance of a type (`<transform>`), and when creating an instance of a type (`<init>`).

In general, `<record>` is used when your modifications require access to both the key and value of the record. For example, if the database key is needed as a factor in an equation, or to identify an element in a dictionary, then `<record>` is the only descriptor in which this type of modification is possible. The `<record>` descriptor is also convenient to use when the number of changes to be made is small, and does not warrant the effort of writing separate `<transform>` or `<init>` descriptors.

The `<transform>` descriptor has a more limited scope than `<record>`. It is used when changes must potentially be made to all instances of a type (regardless of the context in which that type is used) and access to the old value is necessary. The `<transform>` descriptor does not have access to the database key and value, therefore decisions can only be made based on the old and new instances of the type in question.

Finally, the `<init>` descriptor is useful when access to the old instance is not required in order to properly initialize a type. In most cases, this activity could also be performed by a `<transform>` descriptor that simply ignored the old instance, so `<init>` may seem redundant. However, there is one situation where `<init>` is required: when it is necessary to initialize an instance of a type that is introduced by the new Slice definitions. Since there are no instances of this type in the current database, a `<transform>` descriptor for that type would never be executed.

See Also

- Automatic Database Migration
- Freeze Maps
- FreezeScript Transformation XML Reference
- FreezeScript Descriptor Expression Language

Copyright 2016, ZeroC, Inc.

FreezeScript Transformation XML Reference

This page describes the XML elements comprising the FreezeScript transformation descriptors.

On this page:

- <transformdb> Descriptor Element
- <database> Descriptor Element
- <record> Descriptor Element
- <transform> Descriptor Element
- <init> Descriptor Element
- <iterate> Descriptor Element
- <if> Descriptor Element
- <set> Descriptor Element
- <add> Descriptor Element
- <define> Descriptor Element
- <remove> Descriptor Element
- <fail> Descriptor Element
- <delete> Descriptor Element
- <echo> Descriptor Element

`<transformdb>` Descriptor Element

The top-level descriptor in a descriptor file. It requires at least one `<database>` descriptor, and supports any number of `<transform>` and `<init>` child descriptors. This descriptor has no attributes.

`<database>` Descriptor Element

The attributes of this descriptor define the old and new key and value types for the database to be transformed, and optionally the name of the database to which these types apply. It supports any number of child descriptors, but at most one `<record>` descriptor. The `<database>` descriptor also creates a global scope for user-defined symbols.

The attributes supported by the `<database>` descriptor are described in the following table:

Name	Description
`name`	Specifies the name of the database defined by this descriptor. (Optional)
`key`	Specifies the Slice types of the old and new keys. If the types are the same, only one needs to be specified. Otherwise, the types are separated by a comma.
`value`	Specifies the Slice types of the old and new values. If the types are the same, only one needs to be specified. Otherwise, the types are separated by a comma.

As an example, consider the following `<database>` descriptor. In this case, the Freeze map to be transformed currently has key type `int` and value type `::Employee`, and is migrating to a key type of `string`:

XML

```
<database key="int,string" value="::Employee">
```

`<record>` Descriptor Element

Commences the transformation. Child descriptors are executed for each record in the database, providing the user with an opportunity to examine the record's old key and value, and optionally modify the new key and value. Default transformations, as well as `<transform>` and `<init>` descriptors, are executed before the child descriptors. The `<record>` descriptor introduces the following symbols into a local scope: `oldkey`, `newkey`, `oldvalue`, `newvalue`, `facet`. These symbols are accessible to child descriptors, but not to `<transform>` or `<init>` descriptors. The `oldkey` and `oldvalue` symbols are read-only. The `facet` symbol is a string indicating the facet name of the object in the current record, and is only relevant for Freeze evictor databases.

Use caution when modifying database keys to ensure that duplicate keys do not occur. If a duplicate database key is encountered, transformation fails immediately.

Note that database transformation only occurs if a `<record>` descriptor is present.

Copyright 2016, ZeroC, Inc.

<transform> Descriptor Element

Customizes the transformation for all instances of a type in the *new* Slice definitions. The children of this descriptor are executed after the optional default transformation has been performed. Only one <transform> descriptor can be specified for a type, but a <transform> descriptor is not required for every type. The symbols old and new are introduced into a local scope and represent the old and new values, respectively. The old symbol is read-only. The attributes supported by this descriptor are described in the following table:

Name	Description
type	Specifies the Slice type ID for the type's new definition.
default	If false, no default transformation is performed on values of this type. If not specified, the default value is true.
base	This attribute determines whether <transform> descriptors of base class types are executed. If true, the <transform> descriptor of the immediate base class is invoked. If no descriptor is found for the immediate base class, the class hierarchy is searched until a descriptor is found. The execution of any base class descriptors occurs after execution of this descriptor's children. If not specified, the default value is true.
rename	Indicates that a type in the old Slice definitions has been renamed to the new type identified by the type attribute. The value of this attribute is the type ID of the old Slice definition. Specifying this attribute relaxes the strict compatibility rules for enum, struct and class types.

Below is an example of a <transform> descriptor that initializes a new data member:

XML

```
<transform type="::Product">
    <set target="new.salePrice" value="old.listPrice * old.discount"/>
</transform>
```

For class types, transformdb first attempts to locate a <transform> descriptor for the object's most-derived type. If no descriptor is found, transformdb proceeds up the class hierarchy in an attempt to find a descriptor. The base object type, Object, is the root of every class hierarchy and is included in the search for descriptors. It is therefore possible to define a <transform> descriptor for type Object, which will be invoked for every class instance.

Note that <transform> descriptors are executed recursively. For example, consider the following Slice definitions:

Slice

```
struct Inner {
    int sum;
};
struct Outer {
    Inner i;
};
```

When transformdb is performing the default transformation on a value of type Outer, it recursively performs the default transformation on the Inner member, then executes the <transform> descriptor for Inner, and finally executes the <transform> descriptor for Outer. However, if default transformation is disabled for Outer, then no transformation is performed on the Inner member and therefore the <transform> descriptor for Inner is not executed.

<init> Descriptor Element

Defines custom initialization rules for all instances of a type in the new Slice definitions. Child descriptors are executed each time the type is instantiated. The typical use case for this descriptor is for types that have been introduced in the new Slice definitions and whose instances require default values different than what transformdb supplies. The symbol value is introduced into a local scope to represent the instance. The attributes supported by this descriptor are described in the following table:

Copyright 2016, ZeroC, Inc.

Name	Description
`type`	Specifies the Slice type ID of the type's new definition.

Here is a simple example of an `<init>` descriptor:

XML

```
<init type="::Player">
    <set target="value.currency" value="100"/>
</init>
```

Note that, like `<transform>`, `<init>` descriptors are executed recursively. For example, if an `<init>` descriptor is defined for a `struct` t ype, the `<init>` descriptors of the `struct`'s members are executed before the `struct`'s descriptor.

`<iterate>` Descriptor Element

Iterates over a dictionary or sequence, executing child descriptors for each element. The symbol names selected to represent the element information may conflict with existing symbols in the enclosing scope, in which case those outer symbols are not accessible to child descriptors. The attributes supported by this descriptor are described in the following table:

Name	Description
`target`	The sequence or dictionary.
`index`	The symbol name used for the sequence index. If not specified, the default symbol is `i`.
`element`	The symbol name used for the sequence element. If not specified, the default symbol is `elem`.
`key`	The symbol name used for the dictionary key. If not specified, the default symbol is `key`.
`value`	The symbol name used for the dictionary value. If not specified, the default symbol is `value`.

Shown below is an example of an `<iterate>` descriptor that sets the new data member `reviewSalary` to `true` if the employee's salary is greater than $3000:

XML

```
<iterate target="new.employeeMap" key="id" value="emp">
    <if test="emp.salary > 3000">
        <set target="emp.reviewSalary" value="true"/>
    </if>
</iterate>
```

`<if>` Descriptor Element

Conditionally executes child descriptors. The attributes supported by this descriptor are described in the following table:

Name	Description
`test`	A boolean expression.

Child descriptors are executed if the expression in `test` evaluates to true.

`<set>` Descriptor Element

Modifies a value. The `value` and `type` attributes are mutually exclusive. If `target` denotes a dictionary element, that element must already

Copyright 2016, ZeroC, Inc.

exist (i.e., `<set>` cannot be used to add an element to a dictionary). The attributes supported by this descriptor are described in the following table:

Name	Description
`target`	An expression that must select a modifiable value.
`value`	An expression that must evaluate to a value compatible with the target's type.
`type`	If specified, set the target to be an instance of the given Slice class. The value is a type ID from the new Slice definitions. The class must be compatible with the target's type.
`length`	An integer expression representing the desired new length of a sequence. If the new length is less than the current size of the sequence, elements are removed from the end of the sequence. If the new length is greater than the current size, new elements are added to the end of the sequence. If `value` or `type` is also specified, it is used to initialize each new element.
`convert`	If `true`, additional type conversions are supported: between integer and floating point, and between integer and enumeration. Transformation fails immediately if a range error occurs. If not specified, the default value is `false`.

The `<set>` descriptor below modifies a member of a dictionary element:

XML

```
<set target="new.parts['P105J3'].cost" value="new.parts['P105J3'].cost * 1.05"/>
```

This `<set>` descriptor adds an element to a sequence and initializes its value:

XML

```
<set target="new.partsList" length="new.partsList.length + 1" value="'P105J3'"/>
```

As another example, the following `<set>` descriptor changes the value of an enumeration:

XML

```
<set target="new.ingredient" value="::New::Apple"/>
```

Notice in this example that the value refers to a symbol in the new Slice definitions.

<add> Descriptor Element

Adds a new element to a sequence or dictionary. It is legal to add an element while traversing the sequence or dictionary using `<iterate>`, however the traversal order after the addition is undefined. The `key` and `index` attributes are mutually exclusive, as are the `value` and `type` attributes. If neither `value` nor `type` is specified, the new element is initialized with a default value. The attributes supported by this descriptor are described in the following table:

Name	Description
`target`	An expression that must select a modifiable sequence or dictionary.
`key`	An expression that must evaluate to a value compatible with the target dictionary's key type.
`index`	An expression that must evaluate to an integer value representing the insertion position. The new element is inserted before `index`. The value must not exceed the length of the target sequence.

Copyright 2016, ZeroC, Inc.

`value`	An expression that must evaluate to a value compatible with the target dictionary's value type, or the target sequence's element type.
`type`	If specified, set the target value or element to be an instance of the given Slice class. The value is a type ID from the new Slice definitions. The class must be compatible with the target dictionary's value type, or the target sequence's element type.
`convert`	If `true`, additional type conversions are supported: between integer and floating point, and between integer and enumeration. Transformation fails immediately if a range error occurs. If not specified, the default value is `false`.

Below is an example of an `<add>` descriptor that adds a new dictionary element and then initializes its member:

XML

```
<add target="new.parts" key="'P105J4'"/>
<set target="new.parts['P105J4'].cost" value="3.15"/>
```

`<define>` Descriptor Element

Defines a new symbol in the current scope. The attributes supported by this descriptor are described in the following table:

Name	Description
`name`	The name of the new symbol. An error occurs if the name matches an existing symbol in the current scope.
`type`	The name of the symbol's formal Slice type. For user-defined types, the name should be prefixed with `::Old` or `::New` to indicate the source of the type. The prefix can be omitted for primitive types.
`value`	An expression that must evaluate to a value compatible with the symbol's type.
`convert`	If `true`, additional type conversions are supported: between integer and floating point, and between integer and enumeration. Execution fails immediately if a range error occurs. If not specified, the default value is `false`.

Below are two examples of the `<define>` descriptor. The first example defines the symbol `identity` to have type `Ice::Identity`, and proceeds to initialize its members using `<set>`:

XML

```
<define name="identity" type="::New::Ice::Identity"/>
<set target="identity.name" value="steve"/>
<set target="identity.category" value="Admin"/>
```

The second example uses the enumeration we first saw in our discussion of custom database migration to define the symbol `manufacturer` and assign it a default value:

XML

```
<define name="manufacturer" type="::New::BigThree" value="::New::Daimler"/>
```

`<remove>` Descriptor Element

Removes an element from a sequence or dictionary. It is legal to remove an element while traversing a sequence or dictionary using `<iterate>`, however the traversal order after removal is undefined. The attributes supported by this descriptor are described in the following table:

Name	Description

Copyright 2016, ZeroC, Inc.

`target`	An expression that must select a modifiable sequence or dictionary.
`key`	An expression that must evaluate to a value compatible with the key type of the target dictionary.
`index`	An expression that must evaluate to an integer value representing the index of the sequence element to be removed.

`<fail>` Descriptor Element

Causes transformation to fail immediately. If `test` is specified, transformation fails only if the expression evaluates to `true`. The attributes supported by this descriptor are described in the following table:

Name	Description
`message`	A message to display upon transformation failure.
`test`	A boolean expression.

The following `<fail>` descriptor terminates the transformation if a range error is detected:

XML

```
<fail message="range error occurred in ticket count!"
test="old.ticketCount > 32767"/>
```

`<delete>` Descriptor Element

Causes transformation of the current database record to cease, and removes the record from the transformed database. This descriptor has no attributes.

`<echo>` Descriptor Element

Displays values and informational messages. If no attributes are specified, only a newline is printed. The attributes supported by this descriptor are described in the following table:

Name	Description
`message`	A message to display.
`value`	An expression. The value of the expression is displayed in a structured format.

Shown below is an `<echo>` descriptor that uses both `message` and `value` attributes:

XML

```
<if test="old.ticketCount > 32767">
    <echo message="deleting record with invalid ticket count: "
value="old.ticketCount"/>
    <delete/>
</if>
```

See Also

- Custom Database Migration
- Freeze Maps
- Freeze Evictors
- Automatic Database Migration
- FreezeScript Descriptor Expression Language

Copyright 2016, ZeroC, Inc.

Using transformdb

On this page:

- Execution Modes for transformdb
- Using Database Catalogs during Transformation
- Slice Options for transformdb
- Type Options for transformdb
- General Options for transformdb
- Database Arguments for transformdb
- Performing an Automatic Migration
  - Migrating a Single Database
  - Migrating All Databases
- Performing a Migration Analysis
  - Generated File
  - Invocation Modes
- Performing a Custom Migration
- transformdb Usage Strategies
- Transforming Objects
- Using transformdb on an Open Environment

Execution Modes for `transformdb`

The tool operates in one of three modes:

- Automatic migration
- Custom migration
- Analysis

The only difference between automatic and custom migration modes is the source of the transformation descriptors: for automatic migration, `transformdb` internally generates and executes a default set of descriptors, whereas for custom migration the user specifies an external file containing the transformation descriptors to be executed.

In analysis mode, `transformdb` creates a file containing the default transformation descriptors it would have used during automatic migration. You would normally review this file and possibly customize it prior to executing the tool again in its custom migration mode.

Using Database Catalogs during Transformation

Freeze maintains schema information in a catalog for each database environment. If necessary, `transformdb` will use the catalog to determine the names of the databases in the environment, and to determine the key and value types of a particular database. There are two advantages to the tool's use of the catalog:

- Allows `transformdb` to operate on all of the databases in a single invocation
- Eliminates the need for you to specify type information for a database.

For example, you can use automatic migration to transform all of the databases at one time, as shown below:

```
$ transformdb [options] old-env new-env
```

Since we omitted the name of a database to be migrated, `transformdb` uses the catalog in the environment `old-env` to discover all of the databases and their types, generates default transformations for each database, and performs the migration. However, we must still ensure that `transformdb` has loaded the old and new Slice types used by *all* of the databases in the environment.

Slice Options for `transformdb`

The tool supports the standard command-line options common to all Slice processors, with the exception of the include directory (`-I`) option. The options specific to `transformdb` are described below:

- `--old` *`SLICE`*
  `--new` *`SLICE`*
  Loads the old or new Slice definitions contained in the file *`SLICE`*. These options may be specified multiple times if several files must be loaded. However, it is the user's responsibility to ensure that duplicate definitions do not occur (which is possible when two files are loaded that share a common include file). One strategy for avoiding duplicate definitions is to load a single Slice file that contains only `#include` statements for each of the Slice files to be loaded. No duplication is possible in this case if the included files use include guards correctly.

Copyright 2016, ZeroC, Inc.

- `--include-old` *DIR*
  `--include-new` *DIR*
  Adds the directory *DIR* to the set of include paths for the old or new Slice definitions.

Type Options for `transformdb`

In invocation modes for which `transformdb` requires that you define the types used by a database, you must specify one of the following options:

- `--key` *TYPE*`[,`*TYPE*`]`
  `--value` *TYPE*`[,`*TYPE*`]`
  Specifies the Slice type(s) of the database key and value. If the type does not change, then the type only needs to be specified once. Otherwise, the old type is specified first, followed by a comma and the new type. For example, the option `--key int,string` indicates that the database key is migrating from `int` to `string`. On the other hand, the option `--key int,int` indicates that the key type does not change, and could be given simply as `--key int`. Type changes are restricted to those allowed by the compatibility rules, but custom migration provides additional flexibility.

- `-e`
  Indicates that a Freeze evictor database is being migrated. As a convenience, this option automatically sets the database key and value types to those appropriate for the Freeze evictor, and therefore the `--key` and `--value` options are not necessary. Specifically, the key type of a Freeze evictor database is `Ice::Identity`, and the value type is `Freeze::ObjectRecord`. The latter is defined in the Slice file `Freeze/EvictorStorage.ice`; however, this file does not need to be loaded into your old and new Slice definitions.

General Options for `transformdb`

These options may be specified during analysis or migration, as indicated below:

- `-i`
  Requests that `transformdb` ignore type changes that violate the compatibility rules. If this option is not specified, `transformdb` fails immediately if such a violation occurs. With this option, a warning is displayed but `transformdb` continues the requested action. The `-i` option can be specified in analysis or automatic migration modes.

- `-p`
  During migration, this option requests that `transformdb` purge object instances whose type is no longer found in the new Slice definitions.

- `-c`
  Use catastrophic recovery on the old Berkeley DB database environment prior to migration.

- `-w`
  Suppress duplicate warnings during migration. This option is especially useful to minimize diagnostic messages when `transformdb` would otherwise emit the same warning many times, such as when it detects the same issue in every record of a database.

Database Arguments for `transformdb`

In addition to the options described above, `transformdb` accepts as many as three arguments that specify the names of databases and database environments:

- `dbenv`
  The pathname of the old database environment directory.

- `db`
  The name of an existing database file in `dbenv`. `transformdb` never modifies this database.

- `newdbenv`
  The pathname of the database environment directory to contain the transformed database(s). This directory must exist and must not contain an existing database whose name matches a database being migrated.

Performing an Automatic Migration

You can use `transformdb` to automatically migrate one database or all databases in an environment.

Migrating a Single Database

Use the following command line to migrate one database:

Copyright 2016, ZeroC, Inc.

```
$ transformdb [slice-opts] [type-opts] [gen-opts] dbenv db newdbenv
```

If you omit `type-opts`, the tool obtains type information for database `db` from the catalog. For example, consider the following command, which uses automatic migration to transform a database with a key type of `int` and value type of `string` into a database with the same key type and a value type of `long`:

```
$ transformdb --key int --value string,long dbhome data.db newdbhome
```

Note that we did not need to specify the Slice options `--old` or `--new` because our key and value types are primitives. Upon successful completion, the file `newdbhome/data.db` contains our transformed database.

Migrating All Databases

To migrate all databases in the environment, use a command like the one shown below:

```
$ transformdb [slice-opts] [gen-opts] dbenv newdbenv
```

In this invocation mode, you must ensure that `transformdb` has loaded the old and new Slice definitions for all of the types it will encounter among the databases in the environment.

Performing a Migration Analysis

Custom migration is a two-step process: you first write the transformation descriptors, and then execute them to transform a database. To assist you in the process of creating a descriptor file, `transformdb` can generate a default set of transformation descriptors by comparing your old and new Slice definitions. This feature is enabled by specifying the following option:

- `-o` *`FILE`*
  Specifies the descriptor file *`FILE`* to be created during analysis. No migration occurs in this invocation mode.

Generated File

The generated file contains a `<transform>` descriptor for each type that appears in both old and new Slice definitions, and an `<init>` descriptor for types that appear only in the new Slice definitions. In most cases, these descriptors are empty. However, they can contain XML comments describing changes detected by `transformdb` that may require action on your part.

For example, let us revisit the enumeration we defined in our discussion of custom database migration:

Slice

```
enum BigThree { Ford, DaimlerChrysler, GeneralMotors };
```

This enumeration has evolved into the one shown below. In particular, the `DaimlerChrysler` enumerator has been renamed to reflect a corporate name change:

Copyright 2016, ZeroC, Inc.

Slice

```
enum BigThree { Ford, Daimler, GeneralMotors };
```

Next we run `transformdb` in analysis mode:

```
$ transformdb --old old/BigThree.ice --new new/BigThree.ice --key string \
--value ::BigThree -o transform.xml
```

The generated file `transform.xml` contains the following descriptor for the enumeration `BigThree`:

XML

```
<transform type="::BigThree">
    <!-- NOTICE: enumerator `DaimlerChrysler' has been removed -->
</transform>
```

The comment indicates that enumerator `DaimlerChrysler` is no longer present in the new definition, reminding us that we need to add logic in this `<transform>` descriptor to change all occurrences of `DaimlerChrysler` to `Daimler`.

The descriptor file generated by `transformdb` is well-formed and does not require any manual intervention prior to being executed. However, executing an unmodified descriptor file is simply the equivalent of using automatic migration.

Invocation Modes

The sample command line shown in the previous section specified the key and value types of the database explicitly. This invocation mode has the following general form:

```
$ transformdb [slice-opts] [type-opts] [gen-opts] -o FILE
```

Upon successful completion, the generated file contains a `<database>` descriptor that records the type information supplied by `type-opts`, in addition to the `<transform>` and `<init>` descriptors described earlier.

For your convenience, you can omit `type-opts` and allow `transformdb` to obtain type information from the catalog instead:

```
$ transformdb [slice-opts] [gen-opts] -o FILE dbenv
```

In this case, the generated file contains a `<database>` descriptor for each database in the catalog. Note that in this invocation mode, `transformdb` must assume that the names of the database key and value types have not changed, since the only type information available is the catalog in the old database environment. If the tool is unable to locate a new Slice definition for a database's key or value type, it emits a warning message and generates a placeholder value in the output file that you must modify prior to migration.

Performing a Custom Migration

After preparing a descriptor file, either by writing one completely yourself, or modifying one generated by the analysis mode described in the

Copyright 2016, ZeroC, Inc.

previous section, you are ready to migrate a database. One additional option is provided for migration:

- `-f` *`FILE`*
  Execute the transformation descriptors in the file *`FILE`*.

To transform one database, use the following command:

```
$ transformdb [slice-opts] [gen-opts] -f FILE dbenv db newdbenv
```

The tool searches the descriptor file for a `<database>` descriptor whose `name` attribute matches `db`. If no match is found, it searches for a `<database>` descriptor that does not have a `name` attribute.

If you want to transform all databases in the environment, you can omit the database name:

```
$ transformdb [slice-opts] [gen-opts] -f FILE dbenv newdbenv
```

In this case, the descriptor file must contain a `<database>` element for each database in the environment.

Continuing our enumeration example from the analysis discussion above, assume we have modified `transform.xml` to convert the `Chrysler` enumerator, and are now ready to execute the transformation:

```
$ transformdb --old old/BigThree.ice --new new/BigThree.ice -f transform.xml \
dbhome bigthree.db newdbhome
```

`transformdb` Usage Strategies

If it becomes necessary for you to transform a Freeze database, we generally recommend that you attempt to use automatic migration first, unless you already know that custom migration is necessary. Since transformation is a non-destructive process, there is no harm in attempting an automatic migration, and it is a good way to perform a sanity check on your `transformdb` arguments (for example, to ensure that all the necessary Slice files are being loaded), as well as on the database itself. If `transformdb` detects any incompatible type changes, it displays an error message for each incompatible change and terminates without doing any transformation. In this case, you may want to run `transformdb` again with the `-i` option, which ignores incompatible changes and causes transformation to proceed.

Pay careful attention to any warnings that `transformdb` emits, as these may indicate the need for using custom migration. For example, if we had attempted to transform the database containing the `BigThree` enumeration from previous sections using automatic migration, any occurrences of the `Chrysler` enumerator would display the following warning:

```
warning: unable to convert 'Chrysler' to ::BigThree
```

If custom migration appears to be necessary, use analysis to generate a default descriptor file, then review it for `NOTICE` comments and edit as necessary. Liberal use of the `<echo>` descriptor can be beneficial when testing your descriptor file, especially from within the `<record>` descriptor where you can display old and new keys and values.

Transforming Objects

The polymorphic nature of Slice classes can cause problems for database migration. As an example, the Slice parser can ensure that a set of Slice definitions loaded into `transformdb` is complete for all types but classes (and exceptions, but we ignore those because they are not persistent). `transformdb` cannot know that a database may contain instances of a subclass that is derived from one of the loaded classes but whose definition is not loaded. Alternatively, the type of a class instance may have been renamed and cannot be found in the new Slice definitions.

By default, these situations result in immediate transformation failure. However, the `-p` option is a (potentially drastic) way to handle these

Copyright 2016, ZeroC, Inc.

situations: if a class instance has no equivalent in the new Slice definitions and this option is specified, `transformdb` removes the instance any way it can. If the instance appears in a sequence or dictionary element, that element is removed. Otherwise, the database record containing the instance is deleted.

Now, the case of a class type being renamed is handled easily enough using custom migration and the `rename` attribute of the `<transform>` descriptor. However, there are legitimate cases where the destructive nature of the `-p` option can be useful. For example, if a class type has been removed and it is simply easier to start with a database that is guaranteed not to contain any instances of that type, then the `-p` option may simplify the broader migration effort.

This is another situation in which running an automatic migration first can help point out the trouble spots in a potential migration. Using the `-p` option, `transformdb` emits a warning about the missing class type and continues, rather than halting at the first occurrence, enabling you to discover whether you have forgotten to load some Slice definitions, or need to rename a type.

Using `transformdb` on an Open Environment

It is possible to use `transformdb` to migrate databases in an environment that is currently open by another process, but if you are not careful you can easily corrupt the environment and cause the other process to fail. To avoid such problems, you must configure both `transformdb` and the other process to set `Freeze.DbEnv.`*`env-name`*`.DbPrivate=0`. This property has a default value of one, therefore you must explicitly set it to zero. Note that `transformdb` makes no changes to the existing database environment, but it requires exclusive access to the new database environment until transformation is complete.

If you run `transformdb` on an open environment but neglect to set `Freeze.DbEnv.`*`env-name`*`.DbPrivate=0`, you can expect `transformdb` to terminate immediately with an error message stating that the database environment is locked. Before running `transformdb` on an open environment, we strongly recommend that you first verify that the other process was also configured with `Freeze.DbEnv.`*`env-name`*`.DbPrivate=0`.

See Also

- Automatic Database Migration
- Custom Database Migration
- Using the Slice Compilers
- Freeze Catalogs
- Freeze Evictors
- Freeze.*

Copyright 2016, ZeroC, Inc.

Inspecting a Freeze Database

The FreezeScript tool `dumpdb` is used to examine a Freeze database. Its simplest invocation displays every record of the database, but the tool also supports more selective activities. In fact, `dumpdb` supports a scripted mode that shares many of the same XML descriptors as `transformdb`, enabling sophisticated filtering and reporting.

Topics

- Using dumpdb
- FreezeScript Inspection XML Reference

Copyright 2016, ZeroC, Inc.

Using dumpdb

This page describes dumpdb and provides advice on how to best use it.

On this page:

- Overview of Inspection Descriptors
- Inspection Flow of Execution
- Inspection Descriptor Scopes
- Command Line Options for dumpdb
- Database Arguments for dumpdb
- dumpdb Use Cases
  - Dump an Entire Database
  - Dump Selected Records
  - Creating a Sample Descriptor File
  - Executing a Descriptor File
  - Examine the Catalog
- Using dumpdb on an Open Environment

Overview of Inspection Descriptors

dumpdb can read descriptors from an XML file. A dumpdb descriptor file has a well-defined structure. The top-level descriptor in the file is <dumpdb>. A <database> descriptor must be present within <dumpdb> to define the key and value types used by the database. Inside <database>, the <record> descriptor triggers database traversal. Shown below is an example that demonstrates the structure of a minimal descriptor file:

XML

```
<dumpdb>
    <database key="string" value="::Employee">
        <record>
            <echo message="Key: " value="key"/>
            <echo message="Value: " value="value"/>
        </record>
    </database>
</dumpdb>
```

During traversal, type-specific actions are supported by the <dump> descriptor, which is a child of <dumpdb>. One <dump> descriptor may be defined for each type in the Slice definitions. Each time dumpdb encounters an instance of a type, the <dump> descriptor for that type is executed.

The <database>, <record>, and <dump> descriptors may contain general-purpose action descriptors such as <if> and <echo>. These actions resemble statements in programming languages like C++ and Java, in that they are executed in the order of definition and their effects are cumulative. Actions can make use of the FreezeScript expression language.

Although dumpdb descriptors are not allowed to modify the database, they can still define local symbols for scripting purposes. Once a symbol is defined by the <define> descriptor, other descriptors such as <set>, <add>, and <remove> can be used to manipulate the symbol's value.

Inspection Flow of Execution

The descriptors are executed as follows:

- <database> is executed first. Each child descriptor of <database> is executed in the order of definition. If a <record> descriptor is present, database traversal occurs at that point. Any child descriptors of <database> that follow <record> are not executed until traversal completes.
- For each record, dumpdb interprets the key and value, invoking <dump> descriptors for each type it encounters. For example, if the value type of the database is a struct, then dumpdb first attempts to invoke a <dump> descriptor for the structure type, and then recursively interprets the structure's members in the same fashion.

Inspection Descriptor Scopes

Copyright 2016, ZeroC, Inc.

The <database> descriptor creates a global scope, allowing child descriptors of <database> to define symbols that are accessible in any descriptor.

In order for a global symbol to be available to a <dump> descriptor, the symbol must be defined before the <record> descriptor is executed.

Furthermore, certain other descriptors create local scopes that exist only for the duration of the descriptor's execution. For example, the <dump> descriptor creates a local scope and defines the symbol value to represent a value of the specified type. Child descriptors of <dump> can also define new symbols in the local scope, as long as those symbols do not clash with an existing symbol in that scope. It is legal to add a new symbol with the same name as a symbol in an outer scope, but the outer symbol will not be accessible during the descriptor's execution.

The global scope is useful in many situations. For example, suppose you want to track the number of times a certain value was encountered during database traversal. This can be accomplished as shown below:

XML

```
<dumpdb>
    <database key="string" value="::Ice::Identity">
        <define name="categoryCount" type="int" value="0"/>
        <record/>
        <echo message="categoryCount = " value="categoryCount"/>
    </database>
    <dump type="::Ice::Identity">
        <if test="value.category == `Accounting'">
            <set target="categoryCount" value="categoryCount + 1"/>
        </if>
    </dump>
</dumpdb>
```

In this example, the <define> descriptor introduces the symbol categoryCount into the global scope, defining it as type int with an initial value of zero. Next, the <record> descriptor causes traversal to proceed. Each occurrence of the type Ice::Identity causes its <dump> descriptor to be executed, which examines the category member and increases categoryCount if necessary. Finally, after traversal completes, the <echo> descriptor displays the final value of categoryCount.

To reinforce the relationships between descriptors and scopes, consider the diagram in the figure below. Several descriptors are shown, including the symbols they define in their local scopes. In this example, the <iterate> descriptor has a dictionary target and therefore the default symbol for the element value, value, hides the symbol of the same name in the parent <dump> descriptor's scope.

This situation can be avoided by assigning a different symbol name to the element value.

In addition to symbols in the <iterate> scope, child descriptors of <iterate> can also refer to symbols from the <dump> and <database> scopes.

Copyright 2016, ZeroC, Inc.

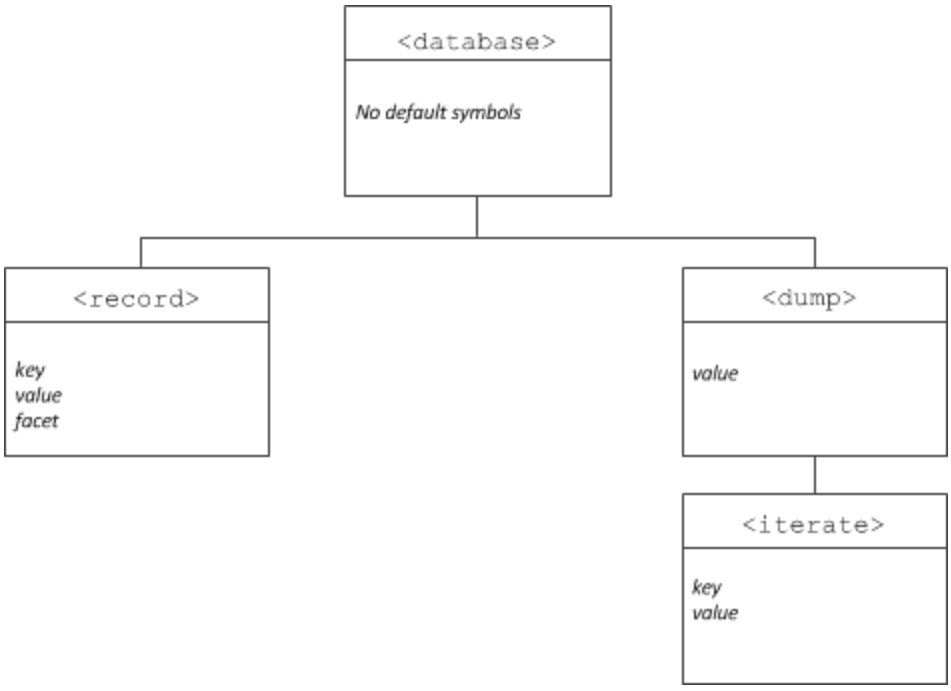

Relationship between descriptors and scopes.

Command Line Options for dumpdb

The tool supports the standard command-line options common to all Slice processors listed. The options specific to `dumpdb` are described below:

- `--load` *`SLICE`*
  Loads the Slice definitions contained in the file *`SLICE`*. This option may be specified multiple times if several files must be loaded. However, it is the user's responsibility to ensure that duplicate definitions do not occur (which is possible when two files are loaded that share a common include file). One strategy for avoiding duplicate definitions is to load a single Slice file that contains only `#include` statements for each of the Slice files to be loaded. No duplication is possible in this case if the included files use include guards correctly.
- `--key` *`TYPE`*
  `--value` *`TYPE`*
  Specifies the Slice type of the database key and value. If these options are not specified, and the `-e` option is not used, `dumpdb` obtains type information from the Freeze catalog.
- `-e`
  Indicates that a Freeze evictor database is being examined. As a convenience, this option automatically sets the database key and value types to those appropriate for the Freeze evictor, and therefore the `--key` and `--value` options are not necessary. Specifically, the key type of a Freeze evictor database is `Ice::Identity`, and the value type is `Freeze::ObjectRecord`. The latter is defined in the Slice file `Freeze/EvictorStorage.ice`, however this file does not need to be explicitly loaded. If this option is not specified, and the `--key` and `--value` options are not used, `dumpdb` obtains type information from the Freeze catalog.
- `-o` *`FILE`*
  Create a file named *`FILE`* containing sample descriptors for the loaded Slice definitions. If type information is not specified, `dumpdb` obtains it from the Freeze catalog. If the `--select` option is used, its expression is included in the sample descriptors. Database traversal does not occur when the `-o` option is used.
- `-f` *`FILE`*
  Execute the descriptors in the file named *`FILE`*. The file's `<database>` descriptor specifies the key and value types; therefore it is not necessary to supply type information.
- `--select` *`EXPR`*
  Only display those records for which the expression *`EXPR`* is true. The expression can refer to the symbols `key` and `value`.
- `-c, --catalog`

Copyright 2016, ZeroC, Inc.

Display information about the databases in an environment, or about a particular database. This option presents the type information contained in the Freeze catalog.

Database Arguments for dumpdb

If dumpdb is invoked to examine a database, it requires two arguments:

- dbenv
  The pathname of the database environment directory.
- db
  The name of the database file. dumpdb opens this database as read-only, and traversal occurs within a transaction.

To display catalog information using the -c option, the database environment directory dbenv is required. If the database file argument db is omitted, dumpdb displays information about every database in the catalog.

dumpdb Use Cases

The command line options support several modes of operation:

- Dump an entire database.
- Dump selected records of a database.
- Emit a sample descriptor file.
- Execute a descriptor file.
- Examine the catalog.

These use cases are described in the following sections.

Dump an Entire Database

The simplest way to examine a database with dumpdb is to dump its entire contents. You must specify the database key and value types, load the necessary Slice definitions, and supply the names of the database environment directory and database file. For example, this command dumps a Freeze map database whose key type is string and value type is Employee:

```
$ dumpdb --key string --value ::Employee --load Employee.ice db emp.db
```

As a convenience, you may omit the key and value types, in which case dumpdb obtains them from the catalog:

```
$ dumpdb --load Employee.ice db emp.db
```

Dump Selected Records

If only certain records are of interest to you, the --select option provides a convenient way to filter the output of dumpdb using an expression. In the following example, we select employees from the accounting department:

```
$ dumpdb --load Employee.ice --select "value.dept == 'Accounting'" db emp.db
```

In cases where the database records contain polymorphic class instances, you must be careful to specify an expression that can be successfully evaluated against all records. For example, dumpdb fails immediately if the expression refers to a data member that does not exist in the class instance. The safest way to write an expression in this case is to check the type of the class instance before referring to any of its data members.

In the example below, we assume that a Freeze evictor database contains instances of various classes in a class hierarchy, and we are only interested in instances of Manager whose employee count is greater than 10:

Copyright 2016, ZeroC, Inc.

```
$ dumpdb -e --load Employee.ice \
--select "value.servant.ice_id == '::Manager' and value.servant.group.length > 10"
\
db emp.db
```

Alternatively, if `Manager` has derived classes, then the expression can be written in a different way so that instances of `Manager` and any of its derived classes are considered:

```
$ dumpdb -e --load Employee.ice \
--select "value.servant.ice_isA('::Manager') and value.servant.group.length > 10" \
db emp.db
```

Creating a Sample Descriptor File

If you require more sophisticated filtering or scripting capabilities, then you must use a descriptor file. The easiest way to get started with a descriptor file is to generate a template using `dumpdb`:

```
$ dumpdb --key string --value ::Employee --load Employee.ice -o dump.xml
```

The output file `dump.xml` is complete and can be executed immediately if desired, but typically the file is used as a starting point for further customization. Again, you may omit the key and value types by specifying the database instead:

```
$ dumpdb --load Employee.ice -o dump.xml db emp.db
```

If the `--select` option is specified, its expression is included in the generated `<record>` descriptor as the value of the `test` attribute in an `<if>` descriptor.

`dumpdb` terminates immediately after generating the output file.

Executing a Descriptor File

Use the `-f` option when you are ready to execute a descriptor file. For example, we can execute the descriptor we generated in the previous section using this command:

```
$ dumpdb -f dump.xml --load Employee.ice db emp.db
```

Examine the Catalog

The `-c` option displays the contents of the database environment's catalog:

```
$ dumpdb -c db
```


Copyright 2016, ZeroC, Inc.

The output indicates whether each database in the environment is associated with an evictor or a map. For maps, the output includes the key and value types.

If you specify the name of a database, `dumpdb` only displays the type information for that database:

```
$ dumpdb -c db emp.db
```

Using dumpdb on an Open Environment

It is possible to use `dumpdb` to migrate databases in an environment that is currently open by another process, but if you are not careful you can easily corrupt the environment and cause the other process to fail. To avoid such problems, you must configure both `dumpdb` and the other process to set `Freeze.DbEnv.`*`env-name`*`.DbPrivate=0`. This property has a default value of one, therefore you must explicitly set it to zero.

If you run `dumpdb` on an open environment but neglect to set `Freeze.DbEnv.`*`env-name`*`.DbPrivate=0`, you can expect `dumpdb` to terminate immediately with an error message stating that the database environment is locked. Before running `dumpdb` on an open environment, we strongly recommend that you first verify that the other process was also configured with `Freeze.DbEnv.`*`env-name`*`.DbPrivate=0`.

See Also

- Using the Slice Compilers
- Freeze Catalogs
- Freeze Evictors
- FreezeScript Descriptor Expression Language
- Freeze.*

Copyright 2016, ZeroC, Inc.

FreezeScript Inspection XML Reference

This page describes the XML elements comprising the FreezeScript inspection descriptors.

On this page:

- <dumpdb> Descriptor Element
- <database> Descriptor Element
- <record> Descriptor Element
- <dump> Descriptor Element
- <iterate> Descriptor Element
- <if> Descriptor Element
- <set> Descriptor Element
- <add> Descriptor Element
- <define> Descriptor Element
- <remove> Descriptor Element
- <fail> Descriptor Element
- <echo> Descriptor Element

<dumpdb> Descriptor Element

The top-level descriptor in a descriptor file. It requires one child descriptor, `<database>`, and supports any number of `<dump>` descriptors. This descriptor has no attributes.

<database> Descriptor Element

The attributes of this descriptor define the key and value types of the database. It supports any number of child descriptors, but at most one `<record>` descriptor. The `<database>` descriptor also creates a global scope for user-defined symbols.

The attributes supported by the `<database>` descriptor are described in the following table:

Name	Description
`key`	Specifies the Slice type of the database key.
`value`	Specifies the Slice type of the database value.

As an example, consider the following `<database>` descriptor. In this case, the Freeze map to be examined has key type `int` and value type `::Employee`:

XML

```
<database key="int" value="::Employee">
```

<record> Descriptor Element

Commences the database traversal. Child descriptors are executed for each record in the database, but after any `<dump>` descriptors are executed. The `<record>` descriptor introduces the read-only symbols `key`, `value` and `facet` into a local scope. These symbols are accessible to child descriptors, but not to `<dump>` descriptors. The `facet` symbol is a string indicating the facet name of the object in the current record, and is only relevant for Freeze evictor databases.

Note that database traversal only occurs if a `<record>` descriptor is present.

<dump> Descriptor Element

Executed for all instances of a Slice type. Only one `<dump>` descriptor can be specified for a type, but a `<dump>` descriptor is not required for every type. The read-only symbol `value` is introduced into a local scope. The attributes supported by this descriptor are described in the following table:

Name	Description
`type`	Specifies the Slice type ID.

Copyright 2016, ZeroC, Inc.

`base`	If `type` denotes a Slice class, this attribute determines whether the <dump> descriptor of the base class is invoked. If `true`, the base class descriptor is invoked after executing the child descriptors. If not specified, the default value is `true`.
`contents`	For `class` and `struct` types, this attribute determines whether descriptors are executed for members of the value. For `sequence` and `dictionary` types, this attribute determines whether descriptors are executed for elements. If not specified, the default value is `true`.

Below is an example of a <dump> descriptor that searches for certain products:

XML

```
<dump type="::Product">
    <if test="value.description.find('scanner') != -1">
        <echo message="Scanner SKU: " value="value.SKU"/>
    </if>
</dump>
```

For class types, `dumpdb` first attempts to locate a <dump> descriptor for the object's most-derived type. If no descriptor is found, `dumpdb` proceeds up the class hierarchy in an attempt to find a descriptor. The base object type, `Object`, is the root of every class hierarchy and is included in the search for descriptors. It is therefore possible to define a <dump> descriptor for type `Object`, which will be invoked for every class instance.

Note that <dump> descriptors are executed recursively. For example, consider the following Slice definitions:

Slice

```
struct Inner {
    int sum;
};
struct Outer {
    Inner i;
};
```

When `dumpdb` is interpreting a value of type `Outer`, it executes the <dump> descriptor for `Outer`, then recursively executes the <dump> descriptor for the `Inner` member, but only if the `contents` attribute of the `Outer` descriptor has the value `true`.

<iterate> Descriptor Element

Iterates over a dictionary or sequence, executing child descriptors for each element. The symbol names selected to represent the element information may conflict with existing symbols in the enclosing scope, in which case those outer symbols are not accessible to child descriptors. The attributes supported by this descriptor are described in the following table:

Name	Description
`target`	The sequence or dictionary.
`index`	The symbol name used for the sequence index. If not specified, the default symbol is `i`.
`element`	The symbol name used for the sequence element. If not specified, the default symbol is `elem`.
`key`	The symbol name used for the dictionary key. If not specified, the default symbol is `key`.
`value`	The symbol name used for the dictionary value. If not specified, the default symbol is `value`.

Shown below is an example of an <iterate> descriptor that displays the name of an employee if the employee's salary is greater than $3000.

Copyright 2016, ZeroC, Inc.

XML

```
<iterate target="value.employeeMap" key="id" value="emp">
    <if test="emp.salary > 3000">
        <echo message="Employee: " value="emp.name"/>
    </if>
</iterate>
```

`<if>` Descriptor Element

Conditionally executes child descriptors. The attributes supported by this descriptor are described in the following table:

Name	Description
`test`	A boolean expression.

Child descriptors are executed if the expression in `test` evaluates to true.

`<set>` Descriptor Element

Modifies a value. The `value` and `type` attributes are mutually exclusive. If `target` denotes a dictionary element, that element must already exist (i.e., `<set>` cannot be used to add an element to a dictionary). The attributes supported by this descriptor are described in the following table:

Name	Description
`target`	An expression that must select a modifiable value.
`value`	An expression that must evaluate to a value compatible with the target's type.
`type`	The Slice type ID of a class to be instantiated. The class must be compatible with the target's type.
`length`	An integer expression representing the desired new length of a sequence. If the new length is less than the current size of the sequence, elements are removed from the end of the sequence. If the new length is greater than the current size, new elements are added to the end of the sequence. If `value` or `type` is also specified, it is used to initialize each new element.
`convert`	If `true`, additional type conversions are supported: between integer and floating point, and between integer and enumeration. Transformation fails immediately if a range error occurs. If not specified, the default value is `false`.

The `<set>` descriptor below modifies a member of a dictionary element:

XML

```
<set target="new.parts['P105J3'].cost" value="new.parts['P105J3'].cost * 1.05"/>
```

This `<set>` descriptor adds an element to a sequence and initializes its value:

XML

```
<set target="new.partsList" length="new.partsList.length + 1" value="'P105J3'"/>
```

`<add>` Descriptor Element

Copyright 2016, ZeroC, Inc.

Adds a new element to a sequence or dictionary. It is legal to add an element while traversing the sequence or dictionary using `<iterate>`, however the traversal order after the addition is undefined. The `key` and `index` attributes are mutually exclusive, as are the `value` and `type` attributes. If neither `value` nor `type` is specified, the new element is initialized with a default value. The attributes supported by this descriptor are described in the following table:

Name	Description
`target`	An expression that must select a modifiable sequence or dictionary.
`key`	An expression that must evaluate to a value compatible with the target dictionary's key type.
`index`	An expression that must evaluate to an integer value representing the insertion position. The new element is inserted before `index`. The value must not exceed the length of the target sequence.
`value`	An expression that must evaluate to a value compatible with the target dictionary's value type, or the target sequence's element type.
`type`	The Slice type ID of a class to be instantiated. The class must be compatible with the target dictionary's value type, or the target sequence's element type.
`convert`	If `true`, additional type conversions are supported: between integer and floating point, and between integer and enumeration. Transformation fails immediately if a range error occurs. If not specified, the default value is `false`.

Below is an example of an `<add>` descriptor that adds a new dictionary element and then initializes its member:

XML

```
<add target="new.parts" key="'P105J4'"/>
<set target="new.parts['P105J4'].cost" value="3.15"/>
```

`<define>` Descriptor Element

Defines a new symbol in the current scope. The attributes supported by this descriptor are described in the following table:

Name	Description
`name`	The name of the new symbol. An error occurs if the name matches an existing symbol in the current scope.
`type`	The name of the symbol's formal Slice type.
`value`	An expression that must evaluate to a value compatible with the symbol's type.
`convert`	If `true`, additional type conversions are supported: between integer and floating point, and between integer and enumeration. Execution fails immediately if a range error occurs. If not specified, the default value is `false`.

Below are two examples of the `<define>` descriptor. The first example defines the symbol `identity` to have type `Ice::Identity`, and proceeds to initialize its members using `<set>`:

XML

```
<define name="identity" type="::Ice::Identity"/>
<set target="identity.name" value="steve"/>
<set target="identity.category" value="Admin"/>
```

The second example uses the enumeration we first saw in our discussion of custom database migration to define the symbol `manufacturer` and assign it a default value:

Copyright 2016, ZeroC, Inc.

XML

```
<define name="manufacturer" type="::BigThree" value="::DaimlerChrysler"/>
```

<remove> Descriptor Element

Removes an element from a sequence or dictionary. It is legal to remove an element while traversing a sequence or dictionary using <iterate>, however the traversal order after removal is undefined. The attributes supported by this descriptor are described in the following table:

Name	Description
target	An expression that must select a modifiable sequence or dictionary.
key	An expression that must evaluate to a value compatible with the key type of the target dictionary.
index	An expression that must evaluate to an integer value representing the index of the sequence element to be removed.

<fail> Descriptor Element

Causes transformation to fail immediately. If test is specified, transformation fails only if the expression evaluates to true. The attributes supported by this descriptor are described in the following table:

Name	Description
message	A message to display upon transformation failure.
test	A boolean expression.

The following <fail> descriptor terminates the transformation if a range error is detected:

XML

```
<fail message="range error occurred in ticket count!"
test="value.ticketCount > 32767"/>
```

<echo> Descriptor Element

Displays values and informational messages. If no attributes are specified, only a newline is printed. The attributes supported by this descriptor are described in the following table:

Name	Description
message	A message to display.
value	An expression. The value of the expression is displayed in a structured format.

Shown below is an <echo> descriptor that uses both message and value attributes:

Copyright 2016, ZeroC, Inc.

XML

```
<if test="value.ticketCount > 32767">
    <echo message="range error occurred in ticket count: "
value="value.ticketCount"/>
</if>
```

See Also

- Freeze Maps
- Freeze Evictors
- Versioning
- Custom Database Migration
- FreezeScript Descriptor Expression Language

Copyright 2016, ZeroC, Inc.

FreezeScript Descriptor Expression Language

An expression language is provided for use in FreezeScript descriptors.

On this page:

- Operators in FreezeScript
- Literals in FreezeScript
- Symbols in FreezeScript
- The nil Keyword in FreezeScript
- Accessing Elements in FreezeScript
- Reserved Keywords in FreezeScript
- Implicit Data Members in FreezeScript
- Calling Functions in FreezeScript
  - String Member Functions
  - Dictionary Member Functions
  - Object Member Functions
  - Global Functions

Operators in FreezeScript

The language supports the usual complement of operators: `and`, `or`, `not`, {+}, `-`, `/`, `*`, `%`, `<`, `>`, `==`, `!=`, `<=`, `>=`, `(`, `)`. Note that the < character must be escaped as `<` in order to comply with XML syntax restrictions.

Literals in FreezeScript

Literal values can be specified for integer, floating point, boolean, and string. The expression language supports the same syntax for literal values as that of Slice, with one exception: string literals must be enclosed in single quotes.

Symbols in FreezeScript

Certain descriptors introduce symbols that can be used in expressions. These symbols must comply with the naming rules for Slice identifiers (i.e., a leading letter followed by zero or more alphanumeric characters). Data members are accessed using dotted notation, such as `value.memberA.memberB`.

Expressions can refer to Slice constants and enumerators using scoped names. In a `transformdb` descriptor, there are two sets of Slice definitions, therefore the expression must indicate which set of definitions it is accessing by prefixing the scoped name with `::Old` or `::New`. For example, the expression `old.fruitMember == ::Old::Pear` evaluates to `true` if the data member `fruitMember` has the enumerated value `Pear`. In `dumpdb`, only one set of Slice definitions is present and therefore the constant or enumerator can be identified without any special prefix.

The `nil` Keyword in FreezeScript

The keyword `nil` represents a nil value of type `Object`. This keyword can be used in expressions to test for a nil object value, and can also be used to set an object value to nil.

Accessing Elements in FreezeScript

Dictionary and sequence elements are accessed using array notation, such as `userMap['joe']` or `stringSeq[5]`. An error occurs if an expression attempts to access a dictionary or sequence element that does not exist. For dictionaries, the recommended practice is to check for the presence of a key before accessing it:

XML

```
<if test="userMap.containsKey('joe') and userMap['joe'].active">
```

(This example shows that you can also call functions in FreezeScript.)

Similarly, expressions involving sequences should check the length of the sequence:

Copyright 2016, ZeroC, Inc.

XML

```
<if test="stringSeq.length > 5 and stringSeq[5] == 'fruit'">
```

The `length` member is an implicit data member.

Reserved Keywords in FreezeScript

The following keywords are reserved: `and`, `or`, `not`, `true`, `false`, `nil`.

Implicit Data Members in FreezeScript

Certain Slice types support implicit data members:

- Dictionary and sequence instances have a member `length` representing the number of elements.
- Object instances have a member `ice_id` denoting the actual type of the object.

Calling Functions in FreezeScript

The expression language supports two forms of function invocation: member functions and global functions. A member function is invoked on a particular data value, whereas global functions are not bound to a data value. For instance, here is an expression that invokes the `find` member function of a `string` value:

```
old.stringValue.find('theSubstring') != -1
```

And here is an example that invokes the global function `stringToIdentity`:

```
stringToIdentity(old.stringValue)
```

If a function takes multiple arguments, the arguments must be separated with commas.

String Member Functions

The `string` data type supports the following member functions:

- `int find(string match[, int start])`
  Returns the index of the substring, or `-1` if not found. A starting position can optionally be supplied.
- `string replace(int start, int len, string str)`
  Replaces a given portion of the string with a new substring, and returns the modified string.
- `string substr(int start[, int len])`
  Returns a substring beginning at the given start position. If the optional length argument is supplied, the substring contains at most `len` characters, otherwise the substring contains the remainder of the string.

Dictionary Member Functions

The `dictionary` data type supports the following member function:

- `bool containsKey(key)`
  Returns `true` if the dictionary contains an element with the given key, or `false` otherwise. The `key` argument must have a value that is compatible with the dictionary's key type.

Copyright 2016, ZeroC, Inc.

Object Member Functions

Object instances support the following member function:

- `bool ice_isA(string id)`
  Returns `true` if the object implements the given interface type, or `false` otherwise. This function cannot be invoked on a nil object.

Global Functions

The following global functions are provided:

- `string generateUUID()`
  Returns a new UUID.
- `string identityToString(Ice::Identity id)`
  Converts an identity into its string representation.
- `string lowercase(string str)`
  Returns a new string converted to lowercase.
- `string proxyToString(Ice::ObjectPrx prx)`
  Returns the string representation of the given proxy.
- `Ice::Identity stringToIdentity(string str)`
  Converts a string into an `Ice::Identity`.
- `Ice::ObjectPrx stringToProxy(string str)`
  Converts a string into a proxy.
- `string typeOf(val)`
  Returns the formal Slice type of the argument.

See Also

- Constants and Literals

Copyright 2016, ZeroC, Inc.

Glacier2

Glacier2 is a lightweight firewall traversal solution for Ice applications.

We present many examples of client/server applications in this manual, most of which assume that the client and server programs are running either on the same host, or on multiple hosts with no network restrictions. We can justify this assumption because this is an instructional text, but a real-world network environment is usually much more complicated: client and server hosts with access to public networks often reside behind protective router-firewalls that not only restrict incoming connections, but also allow the protected networks to run in a private address space using Network Address Translation (NAT). These features, which are practically mandatory in today's hostile network environments, also disrupt the ideal world in which our examples are running.

Topics

- Common Firewall Traversal Issues
- About Glacier2
- How Glacier2 Works
- Getting Started with Glacier2
- Callbacks through Glacier2
- Glacier2 Helper Classes
- Securing a Glacier2 Router
- Glacier2 Session Management
- Dynamic Request Filtering with Glacier2
- Glacier2 Request Buffering
- How Glacier2 uses Request Contexts
- Configuring Glacier2 behind an External Firewall
- Advanced Glacier2 Client Configurations
- IceGrid and Glacier2 Integration
- Glacier2 Metrics

Copyright 2016, ZeroC, Inc.

Common Firewall Traversal Issues

Let's assume that a client and server need to communicate over an untrusted network, and that the client and server hosts reside in private networks behind firewalls:

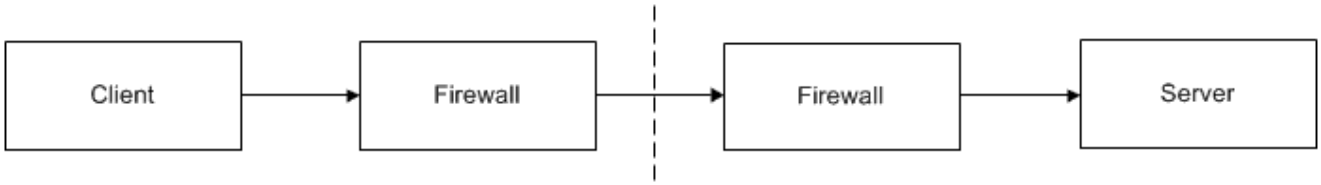

Scenario 1: Client request in a typical network.

Although the diagram looks fairly straightforward, there are several troublesome issues:

- A dedicated port on the server's firewall must be opened and configured to forward messages to the server.
- If the server uses multiple endpoints (e.g., to support both TCP and SSL), then a firewall port must be dedicated to each endpoint.
- The client's proxy must be configured to use the server's "public" endpoint, which is the host name and dedicated port of the firewall.
- If the server returns a proxy as the result of a request, the proxy must not contain the server's private endpoint because that endpoint is inaccessible to the client.

To complicate the scenario even further, the illustration below adds a callback from the server to the client. Callbacks imply that the client is also a server, therefore all of the issues associated with previous illustration now apply to the client as well.

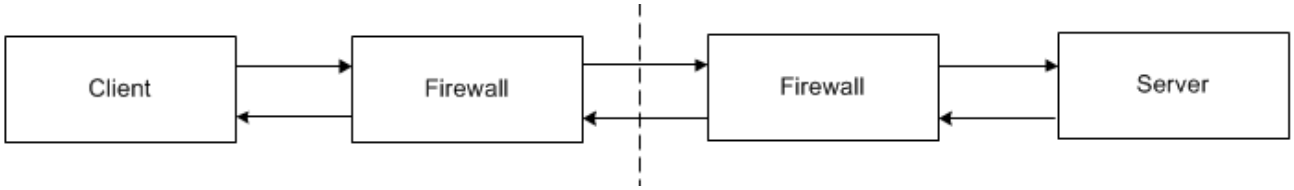

Scenario 2: Callbacks in a typical network.

As if this was not complicated enough already, the illustration below adds multiple clients and servers. Each additional server (including clients requiring callbacks) adds more work for the firewall administrator as more ports are dedicated to forwarding requests.

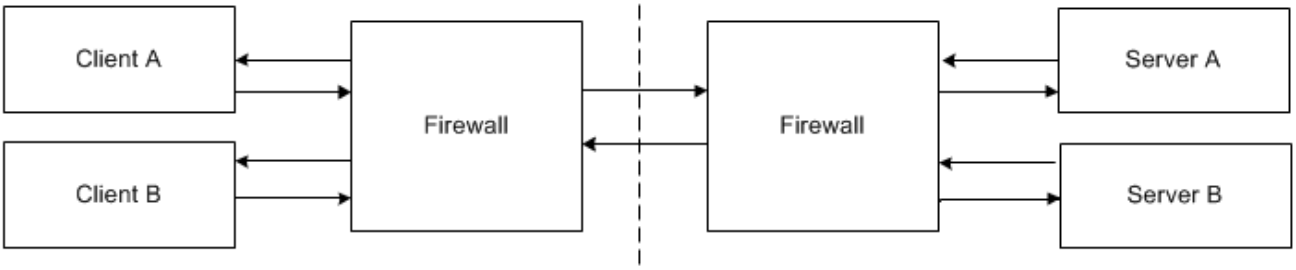

Scenario 3: Multiple clients and servers with callbacks in a typical network.

Clearly, these scenarios do not scale well, and are unnecessarily complex. Fortunately, Ice provides a solution in Glacier2.

See Also

- About Glacier2
- How Glacier2 Works
- Getting Started with Glacier2

Copyright 2016, ZeroC, Inc.

About Glacier2

Glacier2, the router-firewall for Ice applications, addresses common firewall traversal issues with minimal impact on clients or servers (or firewall administrators). In the illustration below, Glacier2 becomes the server firewall for Ice applications. What is not obvious in the diagram, however, is how Glacier2 eliminates much of the complexity of firewall traversal.

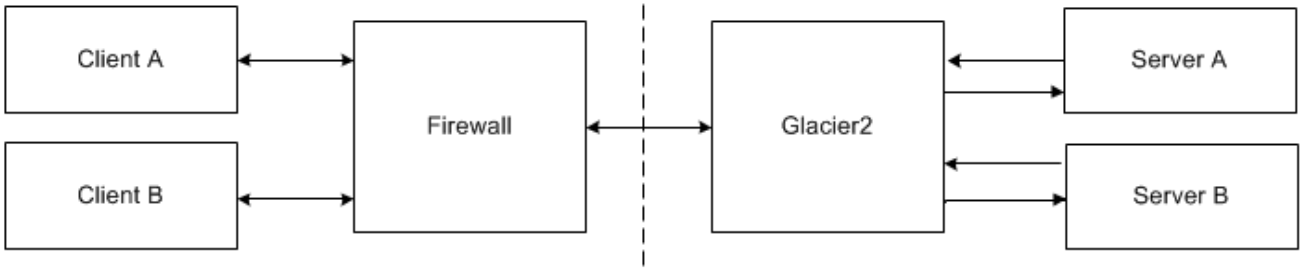

Complex network environments are a fact of life. Unfortunately, the cost of securing an enterprise's network is increased application complexity and administrative overhead. Glacier2 helps to minimize these costs by providing a low-impact, efficient and secure router for Ice applications.

Glacier2 has the following advantages and limitations.

Advantages

- Clients often require only minimal changes to use Glacier2.
- Only one front-end port is necessary to support any number of servers, allowing a Glacier2 router to easily receive connections from a port-forwarding firewall.
- The number of connections to back-end servers is reduced. Glacier2 effectively acts as a connection concentrator, establishing a single connection to each back-end server to forward requests from any number of clients. Similarly, connections from back-end servers to Glacier2 for the purposes of sending callbacks are also concentrated.
- Servers are unaware of Glacier2's presence, and require no modifications whatsoever to use Glacier2. From a server's perspective, Glacier2 is just another local client, therefore servers are no longer required to advertise "public" endpoints in the proxies they create. Furthermore, back-end services such as IceGrid can continue to be used transparently via a Glacier2 router.
- Callbacks through Glacier2 are supported without requiring new connections from servers to clients. In other words, a callback from a server to a client is sent over an existing connection from the client to the server, thereby eliminating the administrative requirements associated with supporting callbacks in the client firewall.
- Glacier2 requires no knowledge of the application's Slice definitions and therefore is very efficient: it routes request and reply messages without unmarshalling the message contents.
- In addition to its primary responsibility of forwarding Ice requests, Glacier2 offers support for user-defined session management and authentication, inactivity timeouts, and request buffering and batching.

Limitations

- Datagram protocols, such as UDP, are not supported.
- Callback objects in a client must use a Glacier2-supplied category in their identities.

See Also

- How Glacier2 Works
- Common Firewall Traversal Issues
- Callbacks through Glacier2
- IceGrid

Copyright 2016, ZeroC, Inc.

How Glacier2 Works

The Ice core supports a generic router facility, represented by the `Ice::Router` interface, that allows a third-party service to "intercept" requests on a properly-configured proxy and deliver them to the intended server. Glacier2 is an implementation of this service, although other implementations are certainly possible.

Glacier2 normally runs on a host in the private network behind a port-forwarding firewall, but it can also operate on a host with access to both public and private networks. In this configuration it follows that Glacier2 must have endpoints on each network.

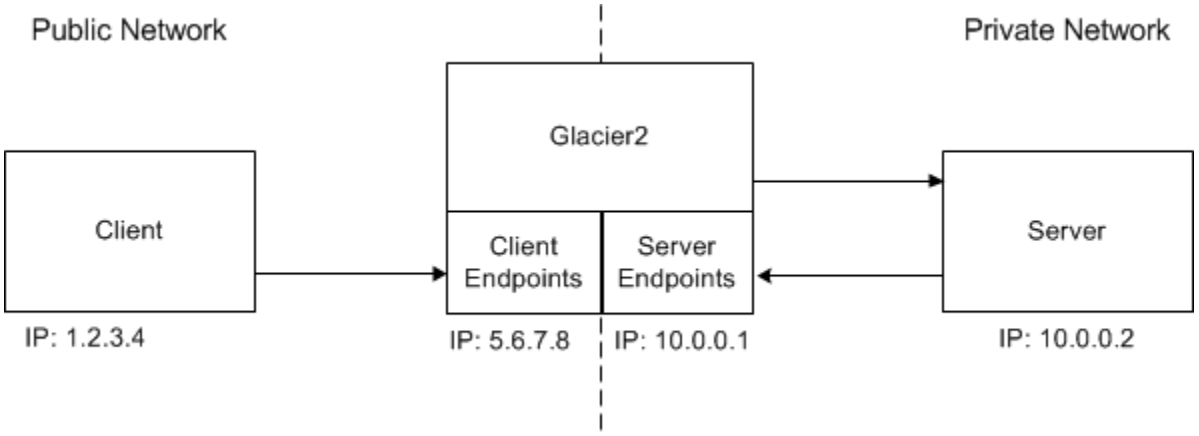

For the sake of example, the router's public address is 5.6.7.8 and its private address is 10.0.0.1.

In the client, proxies must be configured to use Glacier2 as a router. This configuration can be done statically for all proxies created by a communicator, or programmatically for a particular proxy. A proxy configured to use a router is called a *routed proxy*.

When a client invokes an operation on a routed proxy, the client connects to one of Glacier2's client endpoints and sends the request as if Glacier2 is the server. Glacier2 then establishes an outgoing connection to the client's intended server in the private network, forwards the request to that server, and returns the reply (if any) to the client. Glacier2 is essentially acting as a local client on behalf of the remote client.

If a server returns a proxy as the result of an operation, that proxy contains the server's endpoints in the private network, as usual. (Remember, the server is unaware of Glacier2's presence, and therefore assumes that the proxy is usable by the client that requested it.) In the absence of a router, a client would receive an exception if it attempted to use such a proxy. When configured with a router, however, the client ignores the proxy's endpoints and always sends requests to the router's client endpoints instead.

Glacier2's server endpoints, which reside in the private network, are only used when a server makes a callback to a client.

See Also

- Callbacks through Glacier2
- Configuring Glacier2 behind an External Firewall

Copyright 2016, ZeroC, Inc.

Getting Started with Glacier2

On this page:

- Using Glacier2
- Configuring the Router
- Writing a Password File
  - icehashpassword Helper Script
- Starting the Router
- Configuring a Glacier2 Client
- Glacier2 Object Identities
- Creating a Glacier2 Session
- Glacier2 Session Expiration
- Glacier2 Session Destruction

Using Glacier2

Using Glacier2 in a minimal configuration involves the following tasks:

1. Write a configuration file for the router.
2. Write a password file for the router. (Glacier2 also supports other ways to authenticate users.)
3. Decide whether to use the router's internal session manager, or supply your own session manager.
4. Start the router on a host with access to the public and private networks.
5. Modify the client configuration to use the router.
6. Modify the client to create a router session.
7. Ensure that the router session remains active for as long as the client requires it.

For the sake of example, the router's public address is 5.6.7.8 and its private address is 10.0.0.1.

Configuring the Router

The following router configuration properties establish the necessary endpoint and define when a session expires due to inactivity:

```
Glacier2.Client.Endpoints=tcp -h 5.6.7.8 -p 4063
Glacier2.SessionTimeout=60
```

The endpoint defined by `Glacier2.Client.Endpoints` is used by the Ice run time in a client to interact directly with the router. It is also the endpoint where requests from routed proxies are sent. This endpoint is defined on the public network interface because it must be accessible to clients. Furthermore, the endpoint uses a fixed port because clients may be statically configured with a proxy for this endpoint. The port numbers 4063 (for TCP) and 4064 (for SSL) are reserved for Glacier2 by the Internet Assigned Numbers Authority (IANA).

This sample configuration uses TCP as the endpoint protocol, although in most cases, SSL is preferable.

A client must create a session in order to use a Glacier2 router. Our setting for the `Glacier2.SessionTimeout` property causes the router to destroy sessions that have been idle for at least 60 seconds. It is not mandatory to define a timeout, but it is recommended, otherwise session state might accumulate in the router.

Note that this configuration enables the router to forward requests from clients to servers. Additional configuration is necessary to support callbacks from servers to clients.

You must also decide which authentication scheme (or schemes) to use. A file-based mechanism is available, as are more sophisticated strategies.

If clients access a location service via the router, additional router configuration is typically necessary.

Writing a Password File

The router's simplest authentication mechanism uses an access control list in a text file consisting of username and password pairs. Passwords are encoded using the `modular crypt format` (`MCF`).

Copyright 2016, ZeroC, Inc.

The general structure of a MCF encoded password hash is: $*identifier*$*content*, where identifier denotes the scheme used for hashing, and content denotes its contents. Glacier2 supports two types of MCF encoded password hashes:

On Windows and OS X:

- PBKDF2 using SHA-1, SHA-256, or SHA-512 as the digest algorithm.

  PBKDF2 does not have a standard form in the MCF specification. In this case Glacier2 uses the same format as passlib.

  - $pbkdf2-digest$rounds$salt$ for SHA-256 and SHA-512.
  - $pbkdf2$rounds$salt$ for SHA-1.

On Linux:

- Crypt using SHA-256, or SHA-512 as the digest algorithm.

The property Glacier2.CryptPasswords specifies the name of the password file:

```
Glacier2.CryptPasswords=passwords
```

The format of the password file is very simple. Each user name-password pair must reside on a separate line, with whitespace separating the user name from the password. For example, the following password file contains an entry for the user name test:

```
test $5$rounds=110000$5rM9XIDChkgEu.S3$ov7yip4NOi1wymAZmamEv1uKPQRB0WzasoJsWMpRT19
```

icehashpassword Helper Script

You can use the icehashpassword helper script to generate these username-password pairs. This script requires Python and pip to be installed. To install this script run:

```
> pip install zeroc-icehashpassword
```

You can now use the command icehashpassword:

```
> icehashpassword
Password:
$5$rounds=110000$5rM9XIDChkgEu.S3$ov7yip4NOi1wymAZmamEv1uKPQRB0WzasoJsWMpRT19
```

You may also specify several optional parameters:

- -d MESSAGE_DIGEST_ALGORITHM, --digest=MESSAGE_DIGEST_ALGORITHM
- -s SALT_SIZE, --salt=SALT_SIZE
- -r ROUNDS, --rounds=ROUNDS

For example,

Copyright 2016, ZeroC, Inc.

```
> python icehashpassword.py -r 25000 -s 32 -d sha256
Password:
$pbkdf2-sha256$25000$pJSSEuKcs9aaE.J8711LyRlD6H0P4Tyn9J7znpNyLsU$Yx7NNLDfwwLeMbZV84
X2rBYBnvrXvK/TDIQiIGabQIM
```

Note that `icehashpassword` generates PBKDF2 hashes on Windows and OS X, and Crypt hashes on Linux.

This authentication scheme is intended for use in simple applications with a few users. Most applications should install their own custom permissions verifier.

Starting the Router

The router supports the following command-line options:

```
$ glacier2router -h
Usage: glacier2router [options]
Options:
-h, --help     Show this message.
-v, --version  Display the Ice version.
--nowarn       Suppress warnings.
```

The `--nowarn` option prevents the router from displaying warning messages at startup when it is unable to contact a permissions verifier object or a session manager object specified by its configuration.

Additional command line options are supported, including those that allow the router to run as a Windows service or Unix daemon, and Ice includes a utility to help you install the router as a Windows service.

Assuming our configuration properties are stored in a file named `config`, you can start the router with the following command:

```
$ glacier2router --Ice.Config=config
```

Configuring a Glacier2 Client

The following properties configure a client to use a Glacier2 router:

```
Ice.Default.Router=Glacier2/router:tcp -h 5.6.7.8 -p 4063
Ice.RetryIntervals=-1
```

The `Ice.Default.Router` property defines the router proxy. Its endpoints must match those in `Glacier2.Client.Endpoints`.

Setting `Ice.RetryIntervals` to -1 disables automatic retries, which are not useful for proxies configured to use a Glacier2 router.

In versions prior to Ice 3.6, Glacier2 clients were required to disable Active Connection Management in order to prevent the

Copyright 2016, ZeroC, Inc.

unintentional closure of a connection to the router. This is no longer necessary as of Ice 3.6. We show how to correctly configure ACM in our discussion of sessions below.

Glacier2 Object Identities

A Glacier2 router hosts two well-known objects. The default identities of these objects are `Glacier2/router` and `Glacier2/admin`, corresponding to the `Glacier2::Router` and `Glacier2::Admin` interfaces, respectively. If an application requires the use of multiple different (that is, not replicated) routers, it is a good idea to assign unique identities to these objects by configuring the routers with different values of the `Glacier2.InstanceName` property, as shown in the following example:

```
Glacier2.InstanceName=PublicRouter
```

This property changes the category of the object identities, which become `PublicRouter/router` and `PublicRouter/admin`. The client's configuration must also be changed to reflect the new identity:

```
Ice.Default.Router=PublicRouter/router:tcp -h 5.6.7.8 -p 4063
```

One exception to this rule is if you deploy multiple Glacier2 routers as replicas, for example, to gain redundancy or to distribute the message-forwarding load over a number of machines. In that case, all the routers must use the same instance name, and the router clients can use proxies with multiple endpoints, such as:

```
Ice.Default.Router=PublicRouter/router:tcp -h 5.6.7.8 -p 4063:tcp -h 6.10.7.8 -p
4063
```

As of Ice 3.6, a client can discover a router's proxy at run time using the `RouterFinder` interface.

Creating a Glacier2 Session

Session management is provided by the `Glacier2::Router` interface:

Copyright 2016, ZeroC, Inc.

Slice

```
module Glacier2 {
    exception PermissionDeniedException {
        string reason;
    };

    interface Router extends Ice::Router {
        Session* createSession(string userId, string password)
            throws PermissionDeniedException,
                   CannotCreateSessionException;

        Session* createSessionFromSecureConnection()
            throws PermissionDeniedException,
                   CannotCreateSessionException;

        idempotent string getCategoryForClient();

        void refreshSession()
            throws SessionNotExistException;

        void destroySession()
            throws SessionNotExistException;

        idempotent long getSessionTimeout();

        idempotent int getACMTimeout();
    };
};
```

The interface defines two operations for creating sessions: `createSession` and `createSessionFromSecureConnection`. The router requires each client to create a session using one of these operations; only after the session is created will the router forward requests on behalf of the client.

The `createSession` operation expects a user name and password and, depending on the router's configuration, returns either a `Session` proxy or nil. When using the default authentication scheme, the given user name and password must match an entry in the router's password file in order to successfully create a session.

The `createSessionFromSecureConnection` operation does not require a user name and password because it authenticates the client using the credentials associated with the client's SSL connection to the router.

To create a session, the client typically obtains the router proxy from the communicator, downcasts the proxy to the `Glacier2::Router` interface, and invokes one of the `create` operations. The sample code below demonstrates how to do it in C++; the code will look very similar in the other language mappings.

Copyright 2016, ZeroC, Inc.

C++

```
Ice::RouterPrx defaultRouter = communicator->getDefaultRouter();
Glacier2::RouterPrx router = Glacier2::RouterPrx::checkedCast(defaultRouter);
string username = ...;
string password = ...;
Glacier2::SessionPrx session;
try
{
    session = router->createSession(username, password);
}
catch(const Glacier2::PermissionDeniedException& ex)
{
    cout << "permission denied:\n" << ex.reason << endl;
}
catch(const Glacier2::CannotCreateSessionException& ex)
{
    cout << "cannot create session:\n" << ex.reason << endl;
}
```

If the router is configured with a session manager, the `createSession` and `createSessionFromSecureConnection` operations may return a proxy for an object implementing the `Glacier2::Session` interface (or an application-specific derived interface). The client receives a null proxy if no session manager is configured.

A non-nil session proxy returned by a `create` operation must be configured with the router that created it because the session object is only accessible via the router. If the router is configured as the client's default router at the time `createSession` or `createSessionFromSecureConnection` is invoked, as is the case in the example above, then the session proxy is already properly configured and nothing else is required. Otherwise, the client must explicitly configure the session proxy with a router using the `ice_router` proxy method.

If the client wishes to destroy the session explicitly, it must invoke `destroySession` on the router proxy. If a client does not destroy its session, the router destroys it automatically when it expires due to inactivity. A client can obtain the inactivity timeout value by calling `getSessionTimeout` and keep the session alive by periodically calling `refreshSession` if necessary. An easier solution for keeping the session alive is to call `getACMTimeout` and use this value to configure the Active Connection Management behavior of the client's connection to the router:

C++

```
int acmTimeout = router->getACMTimeout();
if (acmTimeout > 0) {
    Ice::ConnectionPtr conn = router->ice_getCachedConnection();
    conn->setACM(acmTimeout, IceUtil::None, Ice::HeartbeatAlways);
}
```

The client calls `getACMTimeout` (new in Ice 3.6) to retrieve the router's server-side ACM timeout and passes this value to an invocation of `setACM` on its connection to the router, thereby ensuring the client and router are using a consistent setting. The client also enables automatic heartbeats so that the connection remains active and prevents the router's server-side ACM from closing the connection.

Alternatively, you could configure the client's ACM with the properties `Ice.ACM.Timeout` and `Ice.ACM.Heartbeat`. Note however that these properties affect all connections created by a client. In general, we recommend configuring the connection directly as shown above.

The `getCategoryForClient` operation is used to implement callbacks over bidirectional connections.

Example

Copyright 2016, ZeroC, Inc.

An example of a Glacier2 client is provided in the directory `demo/Glacier2/callback`.

Glacier2 Session Expiration

A Glacier2 router may be configured to destroy sessions after a period of inactivity. This feature allows the router, as well as a custom session manager, to reclaim resources acquired during the session, but it requires some coordination between the router and its clients.

Ideally you would select a session timeout that is long enough to accommodate the usage patterns of your clients. For example, a session timeout of thirty seconds is a reasonable choice for a client that invokes an operation on a back-end server once every five seconds. However, that timeout could disrupt a different client that has long periods of inactivity, such as when its invocations are prompted by human interaction.

If you cannot predict with certainty the usage patterns of your clients, we recommend configuring the clients so that they actively prevent their sessions from expiring. The simplest solution is to enable the heartbeat feature of Active Connection Management, wherein the Ice run time automatically sends a heartbeat message at regular intervals. The Glacier2 router interprets a heartbeat message as an indication that the client is still active and wishes its session to remain alive. ACM settings can be configured for all connections created by a communicator, or individually for a particular connection. Use care to ensure the client's ACM timeout and the router's session timeout are compatible.

Ice includes helper classes that simplify the task of creating a session.

Note that if a session times out, the next client invocation raises `ConnectionLostException`. To re-establish the session, the client must explicitly re-create it. If the client uses callbacks, it must also re-create the callback adapter and re-register its callback servants.

Glacier2 Session Destruction

A router session is destroyed automatically when the session expires, and when a client explicitly destroys its session. The router also destroys a session if certain connection errors occur while attempting to route a request. These errors are represented by the run-time exceptions `SocketException`, `TimeoutException`, and `ProtocolException`. In other words, if any of these exceptions occur while Glacier2 attempts to establish a connection to the target back-end server, or forward a request to the target back-end server, the router automatically destroys the session.

See Also

- Callbacks through Glacier2
- Securing a Glacier2 Router
- Glacier2 Session Management
- Glacier2.*
- Windows Services
- Active Connection Management
- Automatic Retries
- Glacier2 Helper Classes

Copyright 2016, ZeroC, Inc.

Callbacks through Glacier2

Callbacks from servers to clients are commonly used in distributed applications, often to notify the client about an event such as the completion of a long-running calculation or a change to a database record. Unfortunately, supporting callbacks in a complicated network environment presents its own set of problems. Ice overcomes these obstacles using a Glacier2 router and bidirectional connections.

On this page:

- Bidirectional Connections with Glacier2
- Callbacks and Connection Closure
- Configuring the Router for Callbacks
- Configuring the Client's Object Adapter with a Router
- Callback Object Identities
- Nested Invocations with a Router
- Handling Session Timeouts

Example
The `demo/Glacier2/callback` example illustrates the use of callbacks with Glacier2. The `README` file in the directory provides instructions on running the example, and comments in the configuration file describe the properties in detail.

Bidirectional Connections with Glacier2

While a regular unrouted connection allows requests to flow in only one direction (from client to server), a bidirectional connection enables requests to flow in both directions. This capability is necessary to circumvent the network restrictions that commonly cause firewall traversal issues, namely, client-side firewalls that prevent a server from establishing an independent connection directly to the client. By sending callback requests over the existing connection from the client to the server (more accurately, from the client to the router), we have created a virtual connection back to the client.

This diagram shows the steps involved in making a callback using Glacier2:

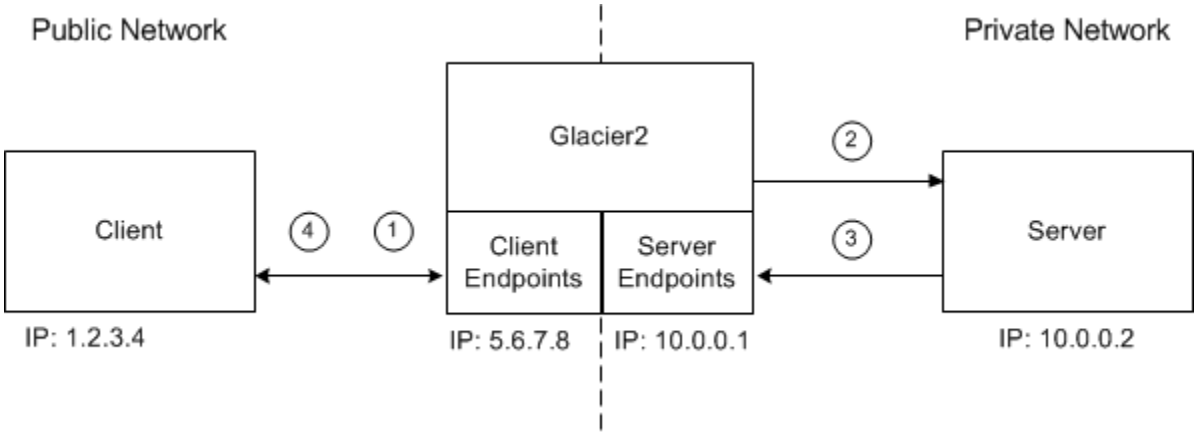

1. The client has a routed proxy for the server and makes an invocation. A connection is established to the router's client endpoint and the request is sent to the router.
2. The router, using information from the client's proxy, establishes a connection to the server and forwards the request. In this example, one of the arguments in the request is a proxy for a callback object in the client.
3. The server makes a callback to the client. For this to succeed, the proxy for the callback object must contain endpoints that are accessible to the server. The only path back to the client is through the router, therefore the proxy contains the router's server endpoints. The server connects to the router and sends the request.
4. The router forwards the callback request to the client using the bidirectional connection established in step 1.

The arrows in the above illustration indicate the flow of requests; notice that two connections are used between the router and the server. Since the server is unaware of the router, it does not use routed proxies, and therefore does not use bidirectional connections.

It is also possible for applications to manually configure bidirectional connections without the use of a router.

Callbacks and Connection Closure

Copyright 2016, ZeroC, Inc.

When a client terminates, it closes its connection to the router. If a server later attempts to make a callback to the client, the attempt fails because the router has no connection to the client over which to forward the request. This situation is no worse than if the server attempted to contact the client directly, which would be prevented by the client's firewall. However, this illustrates the inherent limitation of bidirectional connections: the lifetime of a client's callback proxy is bounded by the lifetime of the client's router session.

Configuring the Router for Callbacks

In order for the router to support callbacks from servers, it needs to have endpoints in the private network.

The configuration file shown below adds the property `Glacier2.Server.Endpoints`:

```
Glacier2.Client.Endpoints=tcp -h 5.6.7.8 -p 4063
Glacier2.Server.Endpoints=tcp -h 10.0.0.1
```

As this example shows, the server endpoint does not require a fixed port.

Configuring the Client's Object Adapter with a Router

A client that receives callbacks is also a server, and therefore must have an object adapter. Typically, an object adapter has endpoints in the local network, but those endpoints are of no use to a server in our restricted network environment. We really want the client's callback proxy to contain the router's server endpoints, and we accomplish that by configuring the client's object adapter with a proxy for the router.

Note that multiple object adapters created by the same communicator cannot use the same router.

We supply the router's proxy by creating the object adapter with `createObjectAdapterWithRouter`, or by defining the object adapter property *`adapter`*`.Router` as shown below:

```
CallbackAdapter.Router=Glacier2/router:tcp -h 5.6.7.8 -p 4063
```

For each object adapter, the Ice run time maintains a list of endpoints that are embedded in proxies created by that adapter. Normally, this list simply contains the local endpoints defined for the object adapter but, when the adapter is configured with a router, the list only contains the router's server endpoints.

An object adapter configured in this way allows the client to receive callback requests via the router. If the client also wants to service requests via local (non-routed) endpoints, the client must create a separate adapter for these requests.

Callback Object Identities

Glacier2 assigns a unique category to each client for use in the identities of the client's callback objects. The client creates proxies that contain this identity category and pass these proxies to back-end servers for use in making callback requests to the client. This category serves two purposes:

1. Upon receipt of a callback request from a back-end server, the router uses the request's category to identify the intended client.
2. The category is sufficiently random that, without knowing the category in advance, it is practically impossible for a misbehaving or malicious back-end server to send callback requests to an arbitrary client.

A client can obtain its assigned category by calling `getCategoryForClient` on the `Router` interface as shown in the example below:

C++

```
Glacier2::RouterPrx router = // ...
string category = router->getCategoryForClient();
```


Copyright 2016, ZeroC, Inc.

Nested Invocations with a Router

If a router client intends to receive callbacks and make nested twoway invocations, it is important that the client be configured correctly. Specifically, you must increase the size of the client thread pool to at least two threads.

Handling Session Timeouts

If the client's session times out, the next invocation raises `ConnectionLostException`. The client can recover from this situation by re-creating the session, re-creating the callback adapter, and adding all the callback servants to the Active Servant Map (ASM) of the re-created adapter.

See Also

- Bidirectional Connections
- Object Adapter Endpoints
- Object Identity
- Getting Started with Glacier2
- The Active Servant Map
- Glacier2.*

Copyright 2016, ZeroC, Inc.

Glacier2 Helper Classes

Ice includes a number of helper classes to help you build robust Glacier2 clients.

- The Glacier2::Application Class
- GUI Helper Classes
  - The SessionFactoryHelper Class
  - The SessionHelper Class
  - The SessionCallback Interface

Compatibility Note
As of Ice 3.6, the Glacier2 helper classes use ACM heartbeats to automatically keep sessions alive. This functionality requires a Glacier2 router that also uses Ice 3.6 or later. Using an earlier version of Glacier2 will require your application to manually keep sessions alive.

The `Glacier2::Application` Class

You may already be familiar with the `Ice::Application` class, which encapsulates some basic Ice functionality such as communicator initialization, communicator destruction, and proper handling of signals and exceptions. The `Glacier2::Application` extends `Ice::Application` to add functionality that is commonly needed by Glacier2 clients:

- Keeps a session alive by periodically sending "heartbeat" requests
- Automatically restarts a session if a failure occurs
- Optionally creates an object adapter for callbacks
- Destroys the session when the application completes

The C++ definition of `Glacier2::Application` is shown below. (The Java and C# versions offer identical functionality so we do not show them here.)

Copyright 2016, ZeroC, Inc.

C++

```
namespace Glacier2 {

class Application : public Ice::Application {
public:
    Application();
    Application(Ice::SignalPolicy policy);

    virtual int runWithSession(int argc, char* argv[]) = 0;

    virtual Glacier2::SessionPrx createSession() = 0;

    virtual void sessionDestroyed();

    static Glacier2::RouterPrx router();

    static Glacier2::SessionPrx session();

    void restart();

    std::string categoryForClient();

    Ice::Identity createCallbackIdentity(const std::string& name);

    Ice::ObjectPrx addWithUUID(const Ice::ObjectPtr& servant);

    Ice::ObjectAdapterPtr objectAdapter();
};

}
```

The class supports the following methods:

- `Application()`
  Instantiating the class using its default constructor has the same semantics as calling `Application(Ice::HandleSignals)`.
- `Application(Ice::SignalPolicy policy)`
  This constructor allows you to indicate whether the class should handle signals. The `SignalPolicy` enumeration contains two enumerators: `HandleSignals` and `NoSignalHandling`. If you specify `HandleSignals`, the class automatically shuts down or destroys its communicator upon receipt of certain signals. Refer to the `Ice::Application` documentation in the relevant language mapping chapter for more information on signal handling.
- `int runWithSession(int argc, char* argv[])`
  This method must be overridden by a subclass and represents the "main loop" of the application. It is called after the communicator has been initialized and the Glacier2 session has been established. The argument vector passed to this method contains the arguments passed to `Application::main` with all Ice-related options removed. The implementation of `runWithSession` must return zero to indicate success and non-zero to indicate failure; the value returned by `runWithSession` becomes the return value of `Application::main`.

  `runWithSession` can call `restart` to restart the session. This destroys the current session, creates a new session (by calling `createSession`), and calls `runWithSession` again. The `Application` base class also restarts the session if `runWithSession` raises (or allows to be raised) any of the following exceptions:
  - `Ice::ConnectionLostException`
  - `Ice::ConnectionRefusedException`
  - `Ice::RequestFailedException`
  - `Ice::TimeoutException`

Copyright 2016, ZeroC, Inc.

- `Ice::UnknownLocalException`

All other exceptions cause the current session to be destroyed without restarting.

- `Glacier2::SessionPrx createSession()`
  This method must be overridden by a subclass to create the application's Glacier2 session. A successful call to `createSession` is followed by a call to `runWithSession`. The application terminates if `createSession` raises (or allows to be raised) an `Ice::LocalException`.
- `void sessionDestroyed()`
  A subclass can optionally override this method to take action when connectivity with the Glacier2 router is lost.
- `Glacier2::RouterPrx router()`
  Returns the proxy for the Glacier2 router.
- `Glacier2::SessionPrx session()`
  Returns the proxy for the current session.
- `void restart()`
  Causes `Application` to destroy the current session, create a new session (by calling `createSession`), and start a new main loop (in `runWithSession`). This method does not return but rather raises a `RestartSessionException` that is trapped by `Application`.
- `std::string categoryForClient()`
  Returns the category to be used in the identities of all of the client's callback objects. Clients must use this category for the router to forward callback requests to the intended client. The method raises `SessionNotExistException` if no session is currently active.
- `Ice::Identity createCallbackIdentity(const std::string& name)`
  Creates a new Ice identity for a callback object with the given identity name.
- `Ice::ObjectPrx addWithUUID(const Ice::ObjectPtr& servant)`
  Adds a servant to the callback object adapter's Active Servant Map using a UUID for the identity name.
- `Ice::ObjectAdapterPtr objectAdapter()`
  Returns the object adapter used for callbacks, creating the object adapter if necessary.

Example
The Ice distribution includes an example in `demo/Glacier2/callback` that shows how to use the `Glacier2::Application` class.

GUI Helper Classes

The "main loop" design imposed by the `Glacier2::Application` class is not suitable for graphical applications, therefore Ice also includes a collection of C++, Java and C# classes that better accommodate the needs of GUI programs:

- `SessionFactoryHelper`
  This class simplifies the task of creating a Glacier2 session. It provides overloaded `connect` methods that support the two authentication styles (user name/password and SSL credentials) accepted by the Glacier2 router, and returns an instance of `SessionHelper` for each new session.
- `SessionHelper`
  This class encapsulates a Glacier2 session and provides much of the same functionality as `Glacier2::Application`. The application must supply an instance of `SessionCallback` when creating the session; `SessionHelper` invokes this callback object when important events occur in the session's lifecycle.
- `SessionCallback`
  An application implements this interface to receive notification about session lifecycle events.

The classes reside in the `Glacier2` namespace or package. We discuss them further in the subsections below.

Example
You can find sample applications that make use of these classes in the `demo/Glacier2/chat` directory of your Ice distribution.

The `SessionFactoryHelper` Class

Copyright 2016, ZeroC, Inc.

The `SessionFactoryHelper` class provides convenience methods for configuring the settings that are commonly used to create a Glacier2 session, such as the router's host and port number. Once the application has completed its configuration, it calls one of the `connect` methods to initialize a communicator, establish a Glacier2 session, and receive a `SessionHelper` object with which it can manage the new session. An application should create a new `SessionFactoryHelper` object for each router instance that it uses.

`SessionFactoryHelper` creates an `Ice.InitializationData` object if the application does not pass one to the `SessionFactoryHelper` constructor. `SessionFactoryHelper` also creates a new property set if necessary, and then sets some configuration properties required by Glacier2 clients. The resulting `InitializationData` object is eventually used in `connect` to initialize a new communicator.

The Java definition of `SessionFactoryHelper` is shown below (the C++ and C# versions are nearly identical and not shown here):

Java

```
package Glacier2;

public class SessionFactoryHelper {

    public SessionFactoryHelper(SessionCallback callback)
        throws Ice.InitializationException;

    public SessionFactoryHelper(Ice.InitializationData initData, SessionCallback
callback)
        throws Ice.InitializationException;

    public SessionFactoryHelper(Ice.Properties properties, SessionCallback
callback)
        throws Ice.InitializationException;

    public void setRouterIdentity(Ice.Identity identity);
    public Ice.Identity getRouterIdentity();

    public void setRouterHost(String hostname);
    public String getRouterHost();

    public void setProtocol(String protocol);
    public String getProtocol();

    public void setTimeout(int timeoutMillisecs);
    public int getTimeout();

    public void setPort(int port);
    public int getPort();

    public Ice.InitializationData getInitializationData();

    public void setConnectContext(java.util.Map<String, String> ctx);

    public void setUseCallbacks(boolean useCallbacks);
    public boolean getUseCallbacks();

    public SessionHelper connect();
    public SessionHelper connect(String username, String password);
}
```

The class supports the following methods:

Copyright 2016, ZeroC, Inc.

- `SessionFactoryHelper(SessionCallback callback)`
  This constructor is useful when your application has no other configuration requirements. The constructor allocates an `InitializationData` object and a new property set. The callback argument must not be null.
- `SessionFactoryHelper(Ice.InitializationData initData, SessionCallback callback)`
  Use this constructor when you want to prepare your own instance of `InitializationData`. The `callback` argument must not be null.
- `SessionFactoryHelper(Ice.Properties properties, SessionCallback callback)`
  This constructor is convenient when you wish to supply an initial set of properties. The callback argument must not be null.
- `void setRouterIdentity(Ice.Identity identity)`
  `Ice.Identity getRouterIdentity()`
  Determines the object identity of the Glacier2 router. If you don't call `setRouterIdentity` to provide an identity, the factory discovers the router's proxy by contacting the `RouterFinder` object using the endpoint information you've configured for the factory (host, port, and so on).

  > The factory ignores this setting if you configure `Ice.Default.Router`.

- `void setRouterHost(String hostname)`
  `String getRouterHost()`
  Determines the host name of the Glacier2 router. The default router host is `localhost`.

  > The factory ignores this setting if you configure `Ice.Default.Router`.

- `void setProtocol(String protocol)`
  `String getProtocol()`
  Determines the Ice protocol used for communications with the Glacier2 router. `protocol` can be `tcp`, `ssl`, `ws` or `wss`; the default protocol is `ssl`.

  > The factory ignores this setting if you configure `Ice.Default.Router`.

- `void setTimeout(int timeoutMillisecs)`
  `int getTimeout()`
  Determines the timeout setting (in milliseconds) for the connection to the Glacier2 router. No timeout is used if the argument is less than or equal to zero. The default timeout is 10,000 ms.

  > The factory ignores this setting if you configure `Ice.Default.Router`.

- `void setPort(int port)`
  `int getPort()`
  Determines the port on which the Glacier2 router is listening. The default is 4063 when the protocol is tcp or ws, and 4064 when the protocol is ssl or wss.

  > The factory ignores this setting if you configure `Ice.Default.Router`.

- `Ice.InitializationData getInitializationData()`
  Returns a reference to the `InitializationData` object that will be used during communicator initialization. If necessary, an application can make modifications to this object prior to calling `connect`.
- `void setConnectContext(java.util.Map<String, String> ctx)`
  Sets the request context to be used when creating a session. This method must be invoked prior to `connect`.
- `void setUseCallbacks(boolean useCallbacks)`
  `boolean getUseCallbacks()`
  Determines whether the session helper automatically creates an object adapter for callback servants. The default behavior is to create the object adapter.
- `SessionHelper connect()`
  Initializes a communicator, creates a Glacier2 session using SSL credentials, and returns a new `SessionHelper` object. The `connected` method is invoked on the session callback if the session was created successfully, otherwise the callback's `connectFaile`

Copyright 2016, ZeroC, Inc.

d method is invoked.

- `SessionHelper connect(String username, String password)`
  Initializes a communicator, creates a Glacier2 session using the given user name and password, and returns a new `SessionHelper` object. The `connected` method is invoked on the session callback if the session was created successfully, otherwise the callback's `connectFailed` method is invoked.

The `SessionHelper` Class

The `SessionHelper` class encapsulates a Glacier2 session and keeps the session alive by periodically "pinging" the router. `SessionHelper` also provides several convenience methods for common session-related actions:

Java

```
package Glacier2;

public class SessionHelper {

    public void destroy();

    public Ice.Communicator communicator();

    public String categoryForClient()
        throws SessionNotExistException;

    public Ice.ObjectPrx addWithUUID(Ice.Object servant)
        throws SessionNotExistException;

    public Glacier2.SessionPrx session()
        throws SessionNotExistException;

    public boolean isConnected();

    public Ice.ObjectAdapter objectAdapter()
        throws SessionNotExistException;
}
```

The class supports the following methods:

- `void destroy()`
  Initiates the destruction of the Glacier2 session, including the communicator created by the `SessionHelper`. This is a non-blocking method that spawns a background thread to perform potentially blocking activities. `SessionHelper` invokes the `disconnected` method on the application's callback object to indicate the completion of the `destroy` request. If the session has already been destroyed when `destroy` is called, the method does nothing.
- `Ice.Communicator communicator()`
  Returns the communicator created by the `SessionHelper`.
- `String categoryForClient()`
  Returns the category that must be used in the identities of all callback objects. Raises `SessionNotExistException` if no session is currently active.
- `Ice.ObjectPrx addWithUUID(Ice.Object servant)`
  Adds a servant to the callback object adapter using a UUID for the identity name. Raises `SessionNotExistException` if no session is currently active.
- `Glacier2.SessionPrx session()`
  Returns a proxy for the Glacier2 session. Raises `SessionNotExistException` if no session is currently active.
- `boolean isConnected()`

Copyright 2016, ZeroC, Inc.

Returns true if the session is currently active, false otherwise.

- `Ice.ObjectAdapter objectAdapter()`
  Returns the callback object adapter. This object adapter is only created if the session factory was configured to do so using `setUseCallbacks`. Raises `SessionNotExistException` if no session is currently active.

The Java implementation of `SessionHelper` installs a JVM shutdown hook that calls `destroy` when the application is about to exit.

The `SessionCallback` Interface

An application must supply an instance of `SessionCallback` when instantiating a `SessionFactoryHelper` object. The callback methods allow the application to receive notification about events in the lifecycle of the session:

Java

```
package Glacier2;

public interface SessionCallback {

    void createdCommunicator(SessionHelper session);

    void connected(SessionHelper session)
        throws SessionNotExistException;

    void disconnected(SessionHelper session);

    void connectFailed(SessionHelper session, Throwable ex);
}
```

The following callback methods are supported:

- `void createdCommunicator(SessionHelper session)`
  Called after successfully initializing a communicator.
- `void connected(SessionHelper session)`
  Called after successfully establishing the Glacier2 session. The method can raise `SessionNotExistException` to force the new session to be destroyed.
- `void disconnected(SessionHelper session)`
  Called after the Glacier2 session is destroyed.
- `void connectFailed(SessionHelper session, Throwable ex)`
  Called if a failure occurred while attempting to establish a Glacier2 session.

See Also

- The Server-Side main Function in C++
- Callbacks through Glacier2
- Glacier2 Session Management

Copyright 2016, ZeroC, Inc.

Securing a Glacier2 Router

As a firewall, a Glacier2 router represents a doorway into a private network, and in most cases that doorway should have a good lock. The obvious first step is to use SSL for the router's client endpoints. This allows you to secure the message traffic and restrict access to clients having the proper credentials. However, the router takes security even further by providing access control and filtering capabilities.

On this page:

- Glacier2 Access Control
  - Password Authentication
  - Certificate Authentication
  - Interaction with a Permissions Verifier
  - Obtaining SSL Credentials for a Router Client
- Request Filtering
  - Address Filters
  - Category Filters
  - Identity Filters
  - Adapter Filters
  - Proxy Filters
  - Client Impact
- Glacier2 Routing Table
- Glacier2 Administrative Interface

Glacier2 Access Control

The authentication capabilities of SSL may not be sufficient for all applications: the certificate validation phase of the SSL handshake verifies that the user is who he says he is, but how do we know that he should be allowed to use the router? Glacier2 addresses this issue through the use of an access control facility that supports two forms of authentication: passwords and certificates. You can configure the router to use whichever authentication method is most appropriate for your application, or you can configure both methods in the same router.

Password Authentication

The router verifies the user name and password arguments to its `createSession` operation before it forwards any requests on behalf of the client. Given that the password is sent "in the clear," it is important to protect these values by using an SSL connection with the router.

There are two ways for the router to verify a user name and password. By default, the router uses a file-based access control list, but you can override this behavior by installing a proxy for an application-defined verifier object. Configuration properties define the password file name or the verifier proxy; if you install a verifier proxy, the password file is ignored. Since we have already discussed the password file, we will focus on the custom verifier interface here.

An application that has special requirements can implement the interface `Glacier2::PermissionsVerifier` to gain programmatic control over access to a router. This can be especially useful in situations where a repository of account information already exists (such as an LDAP directory), in which case duplicating that information in another file would be tedious and error-prone.

The Slice definition for the interface contains just one operation:

Slice

```
module Glacier2 {
    interface PermissionsVerifier {
        idempotent bool checkPermissions(string userId, string password, out string
reason)
            throws PermissionDeniedException;
    };
};
```

The router invokes `checkPermissions` on the verifier object, passing it the user name and password arguments that were given to `creat eSession`. The operation must return true if the arguments are valid, and false otherwise. If the operation returns false, a reason can be

Copyright 2016, ZeroC, Inc.

provided in the output parameter. Starting with Ice 3.5, `checkPermissions` can also throw `Glacier2::PermissionDeniedException`. Glacier2 forwards this exception as-is to the client, which means the verifier can raise a subclass of `PermissionDeniedException` in order to provide more information to the client.

To configure a router with a custom verifier, set the configuration property `Glacier2.PermissionsVerifier` with the proxy for the object.

In situations where authentication is not necessary, such as during development or when running in a trusted environment, you can use Glacier2's built-in "null" permissions verifier. This object accepts any combination of user name and password, and you can enable it with the following property definition:

```
Glacier2.PermissionsVerifier=Glacier2/NullPermissionsVerifier
```

Note that the category of the object's identity (`Glacier2` in this example) must match the value of the property `Glacier2.InstanceName`.

Example
A sample implementation of the `PermissionsVerifier` interface is provided in the `demo/Glacier2/callback` directory.

Certificate Authentication

The `createSessionFromSecureConnection` operation does not require a user name or password because the client's SSL connection to the router already supplies the credentials necessary to sufficiently identify the client, in the form of X.509 certificates.

It is up to you to decide what constitutes sufficient identification. For example, a single certificate could be shared by all clients if there is no need to distinguish between them, or you could generate a unique certificate for each client or a group of clients. Glacier2 does not enforce any particular policy, but simply delegates the decision of whether to accept the client's credentials to an application-defined object that implements the `Glacier2::SSLPermissionsVerifier` interface:

Slice

```
module Glacier2 {
    interface SSLPermissionsVerifier {
        idempotent bool authorize(SSLInfo info, out string reason)
            throws PermissionDeniedException;
    };
};
```

Router clients may only use `createSessionFromSecureConnection` if the router is configured with a proxy for an `SSLPermissionsVerifier` object. The implementation of `authorize` must return true to allow the client to establish a session. To reject the session, `authorize` must return false and may optionally provide a value for `reason`, which is returned to the client as a member of `PermissionDeniedException`. Starting with Ice 3.5, it can also throw `Glacier2::PermissionDeniedException`. Glacier2 forwards this exception as-is to the client, which means the verifier can raise a subclass of `PermissionDeniedException` in order to provide more information to the client.

The verifier examines the members of `SSLInfo` to authenticate a client:

Copyright 2016, ZeroC, Inc.

Slice

```
module Glacier2 {
    struct SSLInfo {
        string remoteHost;
        int remotePort;
        string localHost;
        int localPort;
        string cipher;
        Ice::StringSeq certs;
    };
};
```

The structure includes address information about the remote and local hosts, and a string that describes the ciphersuite negotiated for the SSL connection between the client and the router. These values are generally of interest for logging purposes, whereas the `certs` member supplies the information the verifier needs to make its decision. The client's certificate chain is represented as a sequence of strings that use the Privacy Enhanced Mail (PEM) encoding.

The first element of the sequence corresponds to the client's certificate, followed by its signing certificates. The certificate of the root Certificate Authority (CA) is the last element of the sequence. An empty sequence indicates that the client did not supply a certificate chain.

Although the certificate chain has already been validated by the SSL implementation, a verifier implementation typically needs to examine it in detail before making its decision. As a result, the verifier will need to convert the contents of `certs` into a more usable form. Some Ice platforms, such as Java and .NET, already provide certificate abstractions, and IceSSL supplies its own for C++ users. IceSSL for Java and .NET defines the method `IceSSL.Util.createCertificate`, which accepts a PEM-encoded string and returns an instance of the platform's certificate class. In C++, the class `IceSSL::Certificate` has a constructor that accepts a PEM-encoded string.

In addition to examining certificate attributes such as the distinguished name of the subject and issuer, it is also important that a verifier consider the length of the certificate chain.

To install your verifier, set the `Glacier2.SSLPermissionsVerifier` property with the proxy of your verifier object.

In situations where authentication is not necessary, such as during development or when running in a trusted environment, you can use Glacier2's built-in "null" permissions verifier. This object accepts the credentials of any client, and you can enable it with the following property definition:

```
Glacier2.SSLPermissionsVerifier=Glacier2/NullSSLPermissionsVerifier
```

Note that the category of the object's identity (`Glacier2` in this example) must match the value of the property `Glacier2.InstanceName`.

Interaction with a Permissions Verifier

The router attempts to contact the configured permissions verifiers at startup. If an object is unreachable, the router logs a warning message but continues its normal operation (you can suppress the warning using the `--nowarn` option.) The router does not contact a verifier again until it needs to invoke an operation on the object. For example, when a client asks the router to create a new session, the router makes another attempt to contact the verifier; if the object is still unavailable, the router logs a message and returns `PermissionDeniedExcepti on` to the client.

Obtaining SSL Credentials for a Router Client

Servers that wish to receive information about a client's SSL connection to the router can define the `Glacier2.AddConnectionContext` property. When enabled, the router adds several entries to the request context of each invocation it forwards to a server, providing information such as the client's encoded certificate (if supplied) and addressing details. If the client's connection uses SSL, the router defines the `_con.peerCert` entry in the context. A server can check for the presence of this entry and also extract additional context entries as shown below in this example:

Copyright 2016, ZeroC, Inc.

C++

```
void unlockDoor(string id, const Ice::Current& curr)
{
    Ice::Context::const_iterator i = curr.ctx.find("_con.peerCert");
    if (i != curr.ctx.end()) {
        string certPEM;
        certPEM = i->second;
        cout << "Client address = "
             << curr.ctx["_con.remoteAddress"]
             << ":" << curr.ctx["_con.remotePort"] << endl;
        ...
    }
    ...
}
```

If the client supplied a certificate, the server can decode and examine it using the techniques discussed for IceSSL.

Request Filtering

The Glacier2 router is capable of filtering requests based on a variety of criteria, which helps to ensure that clients do not gain access to unintended objects.

Address Filters

To prevent a client from accessing arbitrary back-end hosts or ports, you can configure a Glacier2 router to validate the address information in each proxy that the client attempts to use. Two properties determine the router's filtering behavior:

- `Glacier2.Filter.Address.Accept`
  An address is accepted if it matches an entry in this property and does not match an entry in `Glacier2.Filter.Address.Reject`.
- `Glacier2.Filter.Address.Reject`
  An address is rejected if it matches an entry in this property.

The value of each property is a list of *address*:*port* pairs separated by spaces, as shown in the example below:

```
Glacier2.Filter.Address.Accept=192.168.1.5:4063 192.168.1.6:4063
```

This configuration allows clients to use only two hosts in the back-end network, and only one port on each host. A client that attempts to use a proxy containing any other host or port receives an `ObjectNotExistException` on its initial request.

You can also use ranges, groups and wildcards when defining your address filters. For example, the following property value shows how to use an address range:

```
Glacier2.Filter.Address.Accept=192.168.1.[5-6]:4063
```

This property is equivalent to the first example, but the range notation allows us to define the filter more concisely. Similarly, we can restate the property using the group notation by separating values with a comma:

Copyright 2016, ZeroC, Inc.

```
Glacier2.Filter.Address.Accept=192.168.1.[5,6]:4063
```

The wildcard notation uses the * character to substitute for a value:

```
Glacier2.Filter.Address.Accept=10.0.*.1:4063
```

The range, group, and wildcard notation is also supported when specifying ports, as shown below:

```
Glacier2.Filter.Address.Accept=192.168.1.[5,6]:[10000-11000]
Glacier2.Filter.Address.Reject=192.168.1.[5,6]:[10500,10501]
```

In this configuration, the router allows clients to access all of the ports in the range `10000` to `11000`, except for the two ports `10500` and `10501`.

At first glance, you might think that the following property definition is pointless because it would prevent clients from accessing any back-end server:

```
Glacier2.Filter.Address.Reject=*
```

In reality, this configuration only prevents clients from accessing servers using direct proxies, that is, proxies that contain endpoints. As a result, the property causes Glacier2 to accept only indirect proxies.

By default, a Glacier2 router forwards requests for any address, which is equivalent to defining the property `Glacier2.Filter.Address.Accept=*`.

Category Filters

The `Ice::Identity` type contains two string members: category and name. You can configure a router with a list of valid identity categories, in which case it only routes requests for objects in those categories. The configuration property `Glacier2.Filter.Category.Accept` supplies the category list:

```
Glacier2.Filter.Category.Accept=cat1 cat2
```

This property does not affect the routing of callback requests from back-end servers to router clients.

By default a Glacier2 router forwards requests for any category.

If a category contains spaces, you can enclose the value in single or double quotes. If a category contains a quote character, it must be escaped with a leading backslash.

Glacier2 can optionally manipulate the category filter automatically. When you set `Glacier2.Filter.Category.AcceptUser` to a value of 1, the router adds the session's user name (for password authentication) or distinguished name (for SSL authentication) to the list of accepted categories. To ensure the uniqueness of your categories, you may prefer setting the property to a value of 2, which causes the router to prepend an underscore to the user name or distinguished name before adding it to the list.

Copyright 2016, ZeroC, Inc.

A session manager can also configure category filters dynamically using Glacier2's `SessionControl` interface.

Identity Filters

The ability to filter on identity categories, as described in the previous section, is a convenient way to limit clients to particular groups of objects. For even stricter control over the identities that clients are allowed to access, you can use the `Glacier2.Filter.Identity.Accept` property. The value of this property is a list of identities, separated by whitespace, representing the *only* objects the router's clients may use.

If an identity contains spaces, you can enclose the value in single or double quotes. If an identity contains a quote character, it must be escaped with a leading backslash.

Clearly, specifying a static list of identities is only practical for a small set of objects. Furthermore, in many applications, the complete set of identities cannot be known in advance, such as when objects are created on a per-session basis and use UUIDs in their identities. For these situations, category-based filtering is generally sufficient. However, a session manager can also use Glacier2's dynamic filtering interface, `SessionControl`, to manage the set of valid identities at run time.

Adapter Filters

Applications often use IceGrid in their back-end network to simplify server administration and take advantage of the benefits offered by indirect proxies. Once you have configured Glacier2 with an appropriate locator proxy, clients can use indirect proxies to refer to objects in IceGrid-managed servers. Indirect proxies come in two forms: one that contains only an identity, and one that contains an identity and an object adapter identifier. You can use the category and identity filters described in previous sections to control identity-only proxies, and you can use the property `Glacier2.Filter.AdapterId.Accept` to enforce restrictions on indirect proxies that use an object adapter identifier.

For example, the following property definition allows a client to use the proxy `factory@WidgetAdapter` but not the proxy `factory@SecretAdapter`:

```
Glacier2.Filter.AdapterId.Accept=WidgetAdapter
```

If an adapter identifier contains spaces, you can enclose the value in single or double quotes. If an adapter identifier contains a quote character, it must be escaped with a leading backslash.

A session manager can also configure this filter dynamically using Glacier2's `SessionControl` interface.

Proxy Filters

The Glacier2 router maintains an internal routing table that contains an entry for each proxy used by a router client; the size of the routing table grows in proportion to the number of clients and their proxy usage. Furthermore, the amount of memory that the routing table consumes is affected by the number of endpoints in each proxy. Glacier2 provides two properties that you can use to limit the size of the routing table and defend against malicious router clients.

The property `Glacier2.RoutingTable.MaxSize` specifies the maximum number of entries allowed in the routing table. If the size of the table exceeds the value of this property, the router evicts older entries on a least-recently-used basis. (Eviction of proxies from the routing table is transparent to router clients.) The default size of the routing table is 1000, but you may need to define a different value depending on the needs of your application. While experimenting with different values, you may find it useful to define the property `Glacier2.Trace.RoutingTable` to see a log of the router's activities with respect to the routing table.

The property `Glacier2.Filter.ProxySizeMax` sets a limit on the size of a stringified proxy. The Ice run time places no limits on the size of proxy components such as identities and host names, but a malicious client could manufacture very large proxies in a denial-of-service attack on a Glacier2 router. By setting this property to a reasonably small value, you can prevent proxies from consuming excessive memory in the router process.

Client Impact

The Glacier2 router immediately terminates a client's session if it attempts to use a proxy that is rejected by an address filter or exceeds the size limit defined by the property `Glacier2.Filter.ProxySizeMax`. The Ice run time in the client responds by raising `ConnectionLostException` to the application.

For category, identity, and adapter identifier filters, the router raises `ObjectNotExistException` if any of the filters rejects a proxy and none of the filters accepts it.

To obtain more information on the router's reasons for terminating a session or rejecting a request, set the following property and examine

Copyright 2016, ZeroC, Inc.

the router's log output:

```
Glacier2.Client.Trace.Reject=1
```

Glacier2 Routing Table

The Glacier2 router maintains an internal routing table for each session. The routing table holds every proxy used by its session. Consequently, the size of the routing table grows in proportion to the number of proxies used by a session. Furthermore, the amount of memory that all of the routing tables consume grows with the number of active sessions.

The property `Glacier2.RoutingTable.MaxSize` allows you to specify an upper limit on the number of entries in the routing table. If the size of the table exceeds the value of this property, the router evicts older entries on a least-recently-used basis. (Eviction of proxies from the routing table is transparent to router clients.) The default size of the routing table is 1000, but you may need to define a different value depending on the needs of your application. While experimenting with different values, you may find it useful to define the property `Glacier2.Trace.RoutingTable` to see a log of the router's activities with respect to the routing table.

The router does not remove entries from a session's routing table except when evicting an old entry to make room for a new one. In particular, an exception that occurs while routing a request for a proxy does *not* cause that proxy to be removed from the routing table. Note however that the routing table is destroyed upon session destruction.

Glacier2 Administrative Interface

Glacier2 supports an administrative interface that allows you to shut down a router programmatically:

Slice

```
module Glacier2 {
    interface Admin {
        idempotent void shutdown();
    };
};
```

To prevent unwanted clients from using the `Admin` interface, the object is only accessible on the endpoints defined by the `Glacier2.Admin.Endpoints` property. This property has no default value, meaning you must define the property in order to make the `Admin` object accessible.

If you decide to define `Glacier2.Admin.Endpoints`, choose your endpoints carefully. We generally recommend the use of endpoints that are accessible only from behind a firewall.

See Also

- Glacier2.*
- IceSSL
- Getting Started with Glacier2
- Callbacks through Glacier2
- Dynamic Request Filtering with Glacier2
- IceGrid and Glacier2 Integration

Copyright 2016, ZeroC, Inc.

Glacier2 Session Management

A Glacier2 router requires a client to create a session and forwards requests on behalf of the client until its session expires. A session expires when it is explicitly destroyed, or when it times out due to inactivity.

You can configure a router to use a custom session manager if your application needs to track the router's session activities. For example, your application may need to acquire resources and initialize the state of back-end services for each new session, and later reclaim those resources when the session expires.

As with the authentication facility, Glacier2 provides two session manager interfaces that an application can implement. The `SessionManager` interface receives notifications about sessions that use password authentication, while the `SSLSessionManager` interface is for sessions authenticated using SSL certificates.

On this page:

- Glacier2 Session Manager Interfaces
- Glacier2 Session Timeouts
- Connection Timeouts on Routed Proxies
- Connection Caching for Session Managers

Glacier2 Session Manager Interfaces

The relevant Slice definitions are shown below:

Slice

```
module Glacier2 {
    exception CannotCreateSessionException {
        string reason;
    };

    interface Session {
        void destroy();
    };

    interface SessionManager {
        Session* create(string userId, SessionControl* control)
            throws CannotCreateSessionException;
    };

    interface SSLSessionManager {
        Session* create(SSLInfo info, SessionControl* control)
            throws CannotCreateSessionException;
    };
};
```

When a client creates a session by invoking `createSession` on the `Router` interface, the router validates the client's user name and password and then calls `SessionManager::create`. Similarly, a call to `createSessionFromSecureConnection` causes the router to invoke `SSLSessionManager::create`. The `SSLInfo` structure provides details about the client's SSL connection. The second argument to the `create` operations is a proxy for a `SessionControl` object, which a session can use to perform dynamic filtering.

The `create` operations must return the proxy of a new `Session` object, or raise `CannotCreateSessionException` and provide an appropriate reason. The `Session` proxy returned by `create` is ultimately returned to the client as the result of `createSession` or `createSessionFromSecureConnection`.

Glacier2 invokes the `destroy` operation on a `Session` proxy when the session expires, giving a custom session manager the opportunity to reclaim resources that were acquired for the session during `create`.

Copyright 2016, ZeroC, Inc.

The `create` operations may be called with information that identifies an existing session. For example, this can occur when a client loses its connection to the router but its previous session has not yet expired (and therefore the router has not yet invoked `destroy` on its `Session` proxy). A session manager implementation must be prepared to handle this situation.

To configure the router with a custom session manager, define the properties `Glacier2.SessionManager` or `Glacier2.SSLSessionManager` with the proxies of the session manager objects. If necessary, you can configure a router with proxies for both types of session managers. If a session manager proxy is not supplied, the call to `createSession` or `createSessionFromSecureConnection` always returns a null proxy.

The router attempts to contact the configured session manager at startup. If the object is unreachable, the router logs a warning message but continues its normal operation (you can suppress the warning using the `--nowarn` option). The router does not contact the session manager again until it needs to invoke an operation on the object. For example, when a client asks the router to create a new session, the router makes another attempt to contact the session manager; if the session manager is still unavailable, the router logs a message and returns `CannotCreateSessionException` to the client.

Example
A sample implementation of the `SessionManager` interface is provided in the `demo/Glacier2/callback` directory.

Glacier2 Session Timeouts

The value of the `Glacier2.SessionTimeout` property specifies the number of seconds a session must be inactive before it expires. This property is not defined by default, which means sessions never expire due to inactivity. If a non-zero value is specified, it's very important that the application chooses a value that does not result in premature session expiration. For example, if it's normal for a client to create a session and then have long periods of inactivity, then a suitably long timeout must be chosen, or the connection must have periodic activity, or timeouts must be disabled altogether. Callbacks from a server to the client also keep a session alive, in which case Glacier2 automatically destroys the session if a failure occurs while forwarding a server callback to the client.

Once a session has expired (or been destroyed for some other reason), the client will no longer be able to send requests via the router and receives a `ConnectionLostException` for each subsequent invocation. The client must explicitly create a new session in order to continue using the router.

In general, we recommend the use of an appropriate session timeout, otherwise resources created for each session will accumulate in the router.

The easiest way for a client to keep a session alive is to enable ACM heartbeats.

Connection Timeouts on Routed Proxies

The timeout for a client's connection to a Glacier2 router works the same as any other connection timeout: in addition to the timeout specified in a proxy's endpoints, several properties can also influence the timeout. For the sake of this discussion, let's ignore the properties and focus on the proxy endpoints.

Before we discuss timeouts, we need to briefly review some Glacier2 concepts. When a Glacier2 client configures a router proxy, normally by setting `Ice.Default.Router`, it causes all proxies to become *routed proxies*. For an invocation on a routed proxy, the Ice run time in the client ignores that proxy's endpoints and instead forwards the invocation to the router for delivery.

There are two kinds of proxies here: the router proxy, and routed proxies. It's important to recognize this distinction in order to understand the following timeout semantics:

- The endpoint timeout in the **router proxy** determines the connection timeout for the connection between the client and the router. Suppose we use the following settings:

```
Ice.Default.Router=Glacier2/router:tcp -h routerHost -p 4063 -t 10000
MyProxy=hello:tcp -h backEndHost -p 8888 -t 5000
```

  This setting configures a ten second timeout for the client's connection to the router. Endpoint timeouts in routed proxies have no effect on this connection.

Copyright 2016, ZeroC, Inc.

- The endpoint timeout in a **routed proxy** affects the router's connection to a back-end server. The Ice run time in the client internally registers routed proxies with the router because the router uses these proxies when it forwards invocations from the client to the desired back-end server. For the routed proxy `MyProxy` in the example above, the router would establish its own connection to port 8888 on `backEndHost` and configure a timeout of five seconds for that connection.

Again, it's critically important that the client's connection to the router remain open for as long as the client requires its session, as closing this connection effectively terminates the session. Since Ice interprets the occurrence of a connection timeout as a hard error and closes the offending connection, select your timeout for the router proxy with care.

As of Ice 3.6, a Glacier2 client can safely use invocation timeouts and active connection management.

For invocations made by Glacier2 to a back-end server, whatever timeout value is set on the *first* proxy that is used to make an invocation applies to all proxies for the same object. This is necessary because Glacier2 adds the proxy to its routing table during the first invocation and, thereafter, reuses that cached proxy for all invocations to the same object identity. Here is an example to illustrate this:

C++

```
// 10-second timeout for the connection to the router.
ObjectPrx router = communicator->stringToProxy("Glacier2/router:tcp -h routerHost
-p 4063 -t 10000");
communicator->setDefaultRouter(RouterPrx::uncheckedCast(router));

// Ping back-end server with 20-second timeout
communicator->stringToProxy("id:tcp -h backEndHost -p 8888 -t 20000")->ice_ping();

// Ping back-end server with 30-second timeout
communicator->stringToProxy("id:tcp -h backEndHost -p 8888 -t 30000")->ice_ping();
```

The first call to `ice_ping`, when forwarded by Glacier2 to the server, establishes a connection with a 20-second timeout. The second call to `ice_ping` reuses the existing connection with a 20-second timeout, even though the proxy specifies a 30-second timeout.

Connection Caching for Session Managers

You can distribute the load among multiple session manager objects by configuring the router with a session manager proxy that contains multiple endpoints. Glacier2 disables connection caching on this proxy so that each invocation on a session manager attempts to use a different endpoint.

This behavior achieves a basic form of load balancing without depending on the replication features provided by IceGrid. Be aware that including an invalid endpoint in your session manager proxy, such as the endpoint of a session manager server that is not currently running, can cause router clients to experience delays during session creation.

If your session managers are in an IceGrid replica group, refer to IceGrid and Glacier2 Integration for more information on the router's caching behavior.

See Also

- Getting Started with Glacier2
- Securing a Glacier2 Router
- Dynamic Request Filtering with Glacier2
- IceGrid and Glacier2 Integration
- Object Adapter Replication
- Glacier2.*

Copyright 2016, ZeroC, Inc.

Dynamic Request Filtering with Glacier2

Glacier2 can be statically configured to filter requests, and also allows a session manager to customize filters for each session at run time via its `SessionControl` interface:

Slice

```
module Glacier2 {
    interface SessionControl {
        StringSet* categories();
        StringSet* adapterIds();
        IdentitySet* identities();
        void destroy();
    };
};
```

The router creates a `SessionControl` object for each client session and supplies a proxy for the object to the session manager `create` operations. Note that the `SessionControl` proxy is null unless the router is configured with server endpoints.

Invoking the `destroy` operation causes the router to destroy the client's session, which eventually results in an invocation of `destroy` on the application-defined `Session` object, if one was provided.

The interface operations `categories`, `adapterIds`, and `identities` return proxies to objects representing the modifiable filters for the session. The router initializes these filters using their respective static configuration properties.

The `SessionControl` object uses a `StringSet` to manage the category and adapter identifier filters:

Slice

```
module Glacier2 {
    interface StringSet {
        idempotent void add(Ice::StringSeq additions);
        idempotent void remove(Ice::StringSeq deletions);
        idempotent Ice::StringSeq get();
    };
};
```

Similarly, the `IdentitySet` interface manages the identity filters:

Slice

```
module Glacier2 {
    interface IdentitySet {
        idempotent void add(Ice::IdentitySeq additions);
        idempotent void remove(Ice::IdentitySeq deletions);
        idempotent Ice::IdentitySeq get();
    };
};
```


Copyright 2016, ZeroC, Inc.

In both interfaces, the `add` operation silently ignores duplicates, and the `remove` operation silently ignores non-existent entries.

Dynamic filtering is often necessary when each session must be restricted to a particular group of objects. Upon session creation, a session manager typically allocates a number of objects in back-end servers for that session to use. To prevent other sessions from accessing these objects (intentionally or not), the session manager can configure the session's filters so that it is only permitted to use the objects that were created for it.

For example, a session manager can retain the `SessionControl` proxy and add a new identity to the `IdentitySet` as each new object is created for the session. A simpler solution is to create a unique category for the session, add it to the session's category filter, and use that category in the identities of all objects accessible to that session. Using a category filter in this way reserves an identity namespace for each session and avoids the need to update the filter for each new object.

To aid in logging and debugging, you can select a category that identifies the client, such as the user name that was supplied during session creation, or an attribute of the client's certificate such as the common name, as long as the selected category is sufficiently unique that it will not conflict with another client's session. You must also ensure that the categories you assign to sessions never match the categories of back-end objects that are not meant to be accessed by router clients. As an example, consider the following session manager implementation:

C++

```
class SessionManagerI : public Glacier2::SessionManager
{
public:

    virtual Glacier2::SessionPrx
    create(const string& username, const Glacier2::SessionControlPrx& ctrl,
           const Ice::Current& curr)
    {
        string category = "_" + username;
        ctrl->categories()->add(category);
        // ...
    }
};
```

This session manager derives a category for the session by prepending an underscore to the user name and then adds this category to the session's filter. As long as our back-end objects do not use a leading underscore in their identity categories, this strategy guarantees that a session's category can never match the category of a back-end object.

For your convenience, Glacier2 already includes support for automatic category filtering.

See Also

- Callbacks through Glacier2
- Securing a Glacier2 Router
- Glacier2 Session Management

Copyright 2016, ZeroC, Inc.

Glacier2 Request Buffering

A Glacier2 router can forward requests in buffered or unbuffered mode. In addition, the buffering mode can be set independently for each direction (client-to-server and server-to-client).

The configuration properties `Glacier2.Client.Buffered` and `Glacier2.Server.Buffered` govern the buffering behavior. The former affects buffering of requests from clients to servers, and the latter affects buffering of requests from servers to clients. If a property is not specified, the default value is `1`, which enables buffering. A property value of `0` selects the unbuffered mode.

The primary difference between the two modes is in the way requests are forwarded:

- Buffered
  The router queues incoming requests and outgoing replies for delivery in a separate thread.
- Unbuffered
  The router forwards requests in the same thread that received the request.

Although unbuffered mode consumes fewer resources than buffered mode, certain features such as request overriding and request batching are available only in buffered mode.

See Also

- How Glacier2 uses Request Contexts

Copyright 2016, ZeroC, Inc.

How Glacier2 uses Request Contexts

The Glacier2 router examines the context of an incoming request for special keys that affect how the router forwards the request. These contexts have the same semantics regardless of whether the request is sent from client to server or from server to client.

On this page:

- The _fwd Context
- The _ovrd Context
- Forwarding Batch Requests
- Context Forwarding

The `_fwd` Context

The `_fwd` context determines the proxy mode that the router uses when forwarding the request. The value associated with the `_fwd` key must be a string containing one or more of the characters shown in the following table:

Value	Mode
d	datagram
D	Batch datagram
o	Oneway
O	Batch oneway
s	Secure
t	Twoway
z	Compress

Legal values for the `_fwd` context key.

These characters match the corresponding stringified proxy options.

For requests whose `_fwd` context specify a batch mode, the forwarding behavior of the router depends on whether it is batching requests.

If the `_fwd` key is not present in a request context, the mode used by the router to forward that request depends on the mode used by the client's proxy as well as the router's own configuration:

- If the client used twoway mode, the router also uses twoway mode.
- If the client sent the request as a oneway or batch oneway, the router's behavior is determined by its configuration properties.

The `_ovrd` Context

In buffered mode, the router allows a new incoming request to override any pending requests that are still in the router's queue, effectively replacing any pending requests with the new request.

For a new request to override a pending request, both requests must meet the following criteria:

- they specify the `_ovrd` key in the request context with the same value
- they are oneway requests
- they are requests on the same object.

This feature is intended to be used by clients that are sending frequent oneway requests in which the most recent request takes precedence. This feature minimizes the number of requests that are forwarded to the server when requests are sent frequently enough that they accumulate in the router's queue before the router has a chance to process them.

Note that the properties `Glacier2.Client.SleepTime` and `Glacier2.Server.SleepTime` can be used to add a delay to the router once it has sent all pending requests. Setting a delay increases the likelihood of overrides actually taking effect. These properties are described in the next section.

Forwarding Batch Requests

Clients can direct the router to forward oneway requests in batches by including the `D` or `O` characters in the `_fwd` context. If the router is

Copyright 2016, ZeroC, Inc.

configured for buffered mode and several such requests accumulate in its queue, the router forwards them together in a batch rather than as individual requests.

In addition, the properties `Glacier2.Client.AlwaysBatch` and `Glacier2.Server.AlwaysBatch` determine whether oneway requests are always batched regardless of the `_fwd` context. The former property affects requests from clients to servers, while the latter affects requests from servers to clients. If a property is defined with a non-zero value, then all requests whose `_fwd` context includes the `o` character or were sent as oneway invocations are treated as if `O` were specified instead, and are batched when possible. Likewise, requests whose `_fwd` context includes the `d` character or were sent as datagram invocations are treated as if `D` were specified. If a property is not defined, the router only batches requests if specifically directed to do so by the `_fwd` context.

The configuration properties `Glacier2.Client.SleepTime` and `Glacier2.Server.SleepTime` can be used to force the router's delivery threads to sleep for the specified number of milliseconds after the router has sent all of its pending requests. (Incoming requests are queued during this period.) The delay is useful to increase the effectiveness of batching because it makes it more likely for additional requests to accumulate in a batch before the batch is sent. If these properties are not defined, or their value is zero, the corresponding thread does not sleep after sending queued requests.

Context Forwarding

The configuration properties `Glacier2.Client.ForwardContext` and `Glacier2.Server.ForwardContext` determine whether the router includes the context when forwarding a request. The former property affects requests from clients to servers, while the latter affects requests from servers to clients. If a property is not defined or has the value zero, the router does not include the context when forwarding requests.

See Also

- Request Contexts
- Proxy and Endpoint Syntax
- Batched Invocations
- Glacier2.*

Copyright 2016, ZeroC, Inc.

Configuring Glacier2 behind an External Firewall

The Glacier2 router requires only one external port to receive connections from clients and therefore can easily coexist with a network firewall device.

For example, consider the network shown in the following illustration:

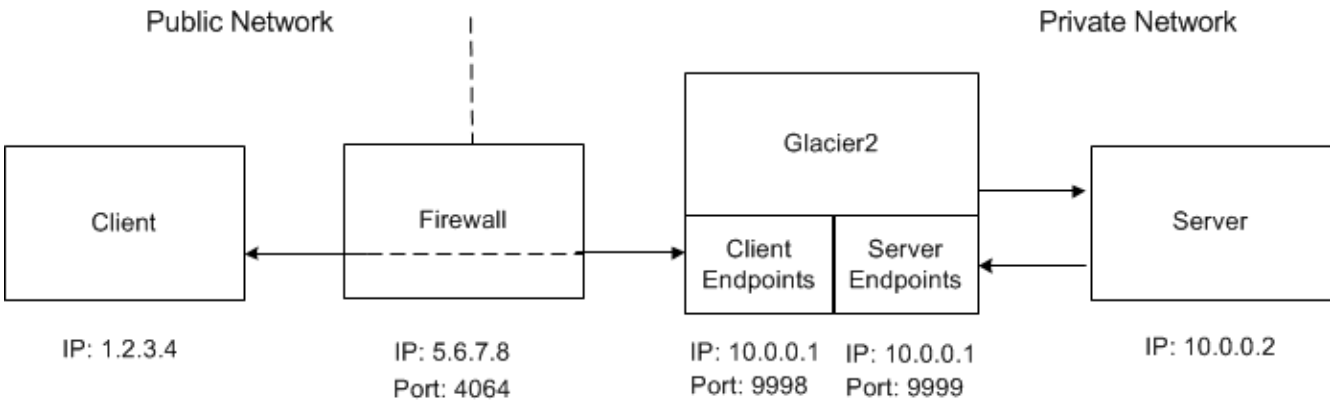

The Glacier2 router in the example above has both of its endpoints in the private network and its host requires only one IP address, unlike the example we showed in the discussion of bidirectional connections in which the Glacier2 host straddled both networks.

We assume that the firewall has been configured to forward connections from port 4064 to the router's client endpoint at port 9998. Meanwhile, the client must be configured to use the firewall's address information in its router proxy, as shown below:

```
Ice.Default.Router=Glacier2/router:ssl -h 5.6.7.8 -p 4064
```

The Glacier2 router configuration for this example requires the following properties:

```
Glacier2.Client.Endpoints=ssl -h 10.0.0.1 -p 9998
Glacier2.Client.PublishedEndpoints=ssl -h 5.6.7.8 -p 4064
Glacier2.Server.Endpoints=tcp -h 10.0.0.1 -p 9999
```

We need to specify published endpoints for the client object adapter because the router is located behind a firewall. Without this property, any proxies that the router creates would use the endpoints specified in `Glacier2.Client.Endpoints`, but of course those endpoints are inaccessible to clients outside the firewall. The `PublishedEndpoints` property forces the Ice run time to use the given endpoints in proxies created by the client object adapter.

Note also that the server endpoint in this example includes a fixed port (`9999`), but a fixed port is not required in the server endpoint for the router to operate properly.

See Also

- Callbacks through Glacier2
- Object Adapter Endpoints
- Glacier2.*

Copyright 2016, ZeroC, Inc.

Advanced Glacier2 Client Configurations

This section details strategies that Glacier2 clients can use to address more advanced requirements.

On this page:

- Callback Strategies with Multiple Object Adapters
- Using Multiple Routers
- Using the RouterFinder Interface

Callback Strategies with Multiple Object Adapters

An application that needs to support callback requests from a router as well as requests from local clients should use multiple object adapters to ensure that proxies created by these object adapters contain the appropriate endpoints. For example, suppose we have the network configuration as shown in the following illustration:

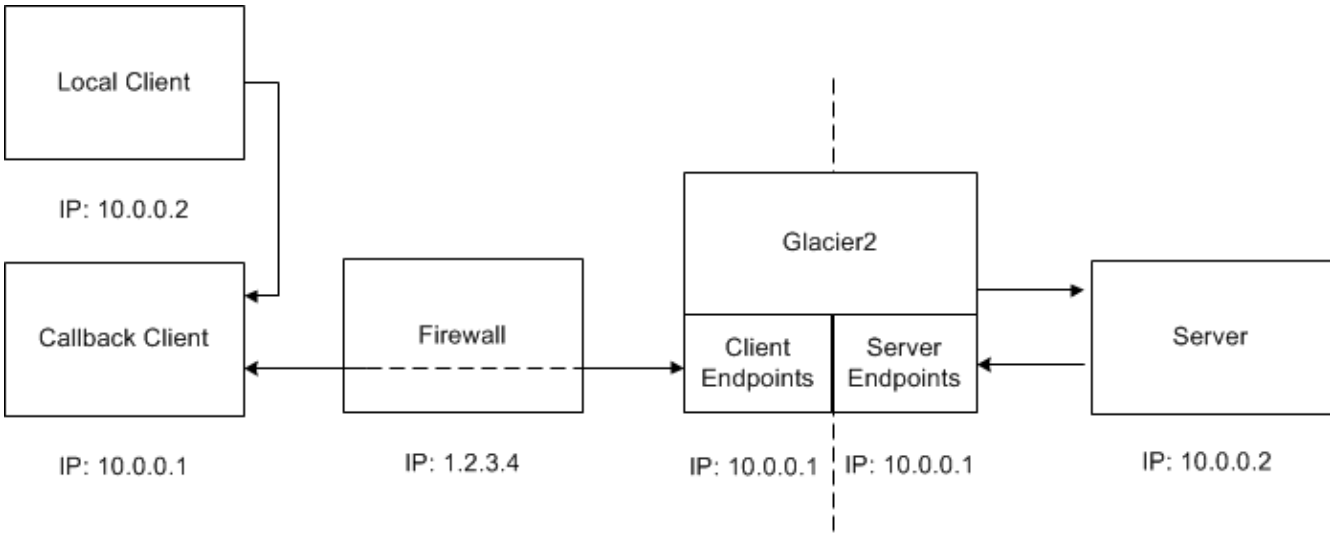

Notice that the two local area networks use the same private network addresses, which is not an unrealistic scenario.

Now, if the callback client were to use a single object adapter for handling both callback requests and local requests, then any proxies created by that object adapter would contain the application's local endpoints as well as the router's server endpoints. As you might imagine, this could cause some subtle problems.

1. When the local client attempts to establish a connection to the callback client via one of these proxies, it might arbitrarily select one of the router's server endpoints to try first. Since the router's server endpoints use addresses in the same network, the local client attempts to make a connection over the local network, with two possible outcomes: the connection attempts to those endpoints fail, in which case they are skipped and the real local endpoints are attempted; or, even worse, one of the endpoints might accidentally be valid in the local network, in which case the local client has just connected to some unintended server.
2. The server may encounter similar problems when attempting to establish a local connection to the router in order to make a callback request.

The solution is to dedicate an object adapter solely to handling callback requests, and another one for servicing local clients. The object adapter dedicated to callback requests must be configured with the router proxy.

Using Multiple Routers

A client is not limited to using only one router at a time: the proxy method `ice_router` allows a client to configure its routed proxies as necessary. With respect to callbacks, a client must create a new callback object adapter for each router that can forward callback requests to the client. A client must also be aware of the object identities in use by the routers.

Using the `RouterFinder` Interface

A router's identity can be changed by setting the `Glacier2.InstanceName` property, which affects the category portion of the identity. The default identity of a Glacier2 router is `Glacier2/Router`, but we can change it to `Production/Router` with the following setting:

Copyright 2016, ZeroC, Inc.

```
Glacier2.InstanceName=Production
```

A client could configure its corresponding router proxy as follows:

```
Ice.Default.Router=Production/Glacier2:tcp -p 4063 -h prodhost
```

In most cases the client can statically configure the router's proxy as we've shown here. For clients that need to discover a router's proxy at run time, Ice also requires router implementations to support the `RouterFinder` interface:

Slice

```
module Ice {
interface RouterFinder {
    Router* getRouter();
};
};
```

An object supporting this interface must be available with the identity `Ice/RouterFinder`. By knowing the host and port of a router's client endpoints, a client can discover the router's proxy with a call to `getRouter`:

C++

```
Ice::ObjectPrx prx = communicator->stringToProxy("Ice/RouterFinder:tcp -p 4063 -h
prodhost");
Ice::RouterFinderPrx finder = Ice::RouterFinderPrx::checkedCast(prx);
Ice::RouterPrx router = finder->getRouter();
communicator->setDefaultRouter(router);
```

The `RouterFinder` interface was added in Ice 3.6.

See Also

- Getting Started with Glacier2
- Callbacks through Glacier2
- Proxy Methods

Copyright 2016, ZeroC, Inc.

IceGrid and Glacier2 Integration

IceGrid is a server activation and location service. This section describes the ways in which you can integrate Glacier2 and IceGrid.

On this page:

- Configuring Router Clients for IceGrid
- Using Replicated Session Managers

Configuring Router Clients for IceGrid

It is not uncommon for a Glacier2 client to require access to a locator service such as IceGrid. In the absence of Glacier2, a locator client would typically define the property `Ice.Default.Locator` with a stringified proxy for the locator service. However, when that locator service is accessed via a Glacier2 router, the configuration requirements are slightly different. It is no longer necessary for the client's configuration to include `Ice.Default.Locator`; this property must be defined in the router's configuration instead.

For example, consider the following network architecture:

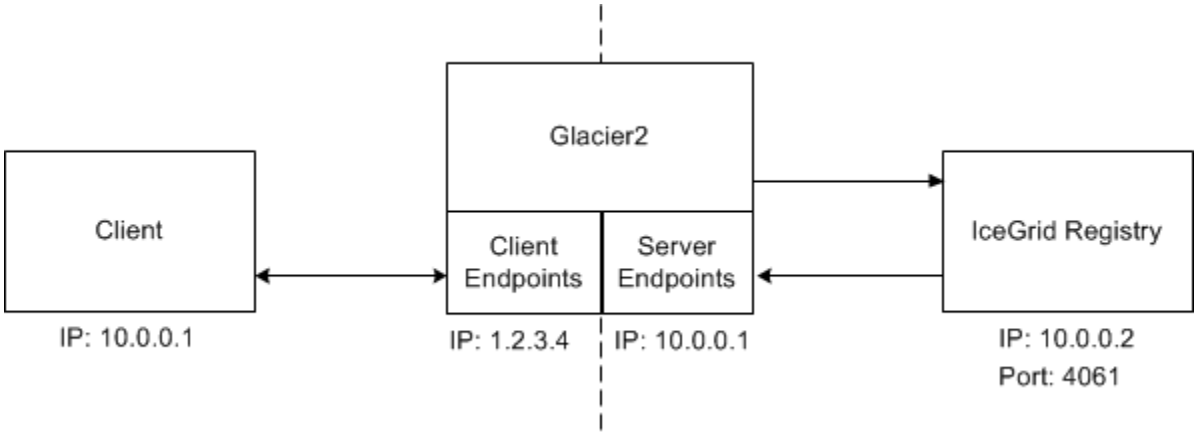

In this case the Glacier2 router's configuration must include the property shown below:

```
Ice.Default.Locator=IceGrid/Locator:tcp -h 10.0.0.2 -p 4061
```

Using Replicated Session Managers

An IceGrid application might want to use replication to increase the availability of Glacier2 session managers. When you configure an indirect proxy for a session manager (and configure Glacier2 with a locator proxy), the Ice run time in the router queries the locator to obtain a proxy for a session manager replica.

By default, this proxy is cached for 10 minutes, meaning the router uses the same session manager proxy to create sessions for a 10-minute period, after which it queries the locator again. If you want to distribute the session-creation load among the session manager replicas more evenly, you can decrease the locator cache timeout using configuration properties.

For example, the following settings use a timeout of 30 seconds:

```
Glacier2.SessionManager.LocatorCacheTimeout=30
Glacier2.SSLSessionManager.LocatorCacheTimeout=30
```

As you can see, timeouts are specified individually for the `SessionManager` and `SSLSessionManager` proxies. You can also disable caching completely by using a value of 0, in which case the router queries the locator before every invocation on a session manager. See the discussion of session management for more details.

Copyright 2016, ZeroC, Inc.

See Also

- IceGrid
- Glacier2 Session Management
- Getting Started with IceGrid
- Object Adapter Replication

Copyright 2016, ZeroC, Inc.

Glacier2 Metrics

You can monitor Glacier2 using the Administrative Facility and the Metrics Facet. Glacier2 provides a metrics class to monitor session related metrics. The Glacier2 session metrics class is defined in `Glacier2/Metrics.ice` and is shown below.

Slice

```
namespace IceMX {
class SessionMetrics extends Metrics {
    int forwardedClient = 0;
    int forwardedServer = 0;
    int routingTableSize = 0;
    int queuedClient = 0;
    int queuedServer = 0;
    int overriddenClient = 0;
    int overriddenServer = 0;
};
};
```

Glacier2 records session metrics in the `Session` metrics map, the metrics objects contained in this map are instances of the `IceMX::SessionMetrics` class show above. To configure a metrics view to record Glacier2 session metrics you can use metrics properties with the `IceMX.Metrics.`*`view-name`*`.Map.Session` prefix, for example:

- `IceMX.Metrics.SessionView.Map.Session.GroupBy=id` to configure a view containing one metrics object per session.
- `IceMX.Metrics.SessionView.Map.Session.GroupBy=none` to configure a view containing a single metrics object with metrics for all the sessions.

You can use the following attributes when configuring the Glacier2 Session metrics map:

Name	Description
id	A unique identifier to identify the session. This corresponds to the user Id for sessions created with a username/password and to the subject DN for sessions created from SSL.
parent	The Glacier2 router instance name.
none	The empty string.
endpoint	The stringified endpoint.
endpointType	The endpoint numerical type as defined in `Ice/Endpoint.ice`.
endpointIsDatagram	A boolean indicating if the endpoint is a datagram endpoint.
endpointIsSecture	A boolean indicating if the endpoint is secure.
endpointTimeout	The endpoint timeout.
endpointCompress	A boolean indicating if the endpoint requires compression.
endpointHost	The endpoint host.
endpointPort	The endpoint port.
connection	The connection description.
incoming	A boolean indicating if the connection is a server or client connection.
adapterName	If the connection is a server connection, adapterName will return the name of the adapter which created the connection, it will contain the empty string otherwise.

Copyright 2016, ZeroC, Inc.

connectionId	The ID of the connection if one is set, the empty string otherwise.
localAddress	The connection's local address.
localPort	The connection's local port.
remoteAddress	The connection's remote address.
remotePort	The connection's remote port.
mcastAddress	The connection's multicast address.
mcastPort	The connection's multicast port.
state	The state of the connection.

The connection and endpoint attributes are for the connection tied to the Glacier2 session.

See Also

- Administrative Facility
- The Metrics Facet
- IceMX.Metrics.*

Copyright 2016, ZeroC, Inc.

IceGrid

IceGrid is a location and activation service for Ice applications. For the purposes of this discussion, we can loosely define *grid computing* as the use of a network of relatively inexpensive computers to perform the computational tasks that once required costly "big iron." Developers familiar with distributed computing technologies may not consider the notion of grid computing to be particularly revolutionary; after all, distributed applications have been running on networks for years, and the definition of grid computing is sufficiently vague that practically any server environment could be considered a "grid."

One possible grid configuration is a homogeneous collection of computers running identical programs. Each computer in the grid is essentially a clone of the others, and all are equally capable of handling a task. As a developer, you need to write the code that runs on the grid computers, and Ice is ideally suited as the infrastructure that enables the components of a grid application to communicate with one another. However, writing the application code is just the first piece of the puzzle. Many other challenges remain:

- How do I install and update this application on all of the computers in the grid?
- How do I keep track of the servers running on the grid?
- How do I distribute the load across all the computers?
- How do I migrate a server from one computer to another one?
- How can I quickly add a new computer to the grid?

Of course, these are issues faced by most distributed applications. As you learn more about IceGrid's capabilities, you will discover that it offers solutions to these challenges. To get you started, we have summarized IceGrid's feature set below:

- Location service
  As an implementation of an Ice location service, IceGrid enables clients to bind indirectly to their servers, making applications more flexible and resilient to changing requirements. (IceDiscovery is an alternate, lightweight location service implementation for applications that don't require IceGrid's additional features.)

- On-demand server activation
  Starting an Ice server process is called *server activation*. IceGrid can be given responsibility for activating a server on demand, that is, when a client attempts to access an object hosted by the server. Activation usually occurs as a side effect of indirect binding, and is completely transparent to the client.

- Application distribution
  IceGrid provides a convenient way to distribute your application to a set of computers, without the need for a shared file system or complicated scripts. Simply configure an IcePatch2 server and let IceGrid download the necessary files and keep them synchronized.

- Replication and load balancing
  IceGrid supports replication by grouping the object adapters of several servers into a single virtual object adapter. During indirect binding, a client can be bound to an endpoint of any of these adapters. Furthermore, IceGrid monitors the load on each computer and can use that information to decide which of the endpoints to return to a client.

- Sessions and resource allocation
  An IceGrid client establishes a session in order to allocate a resource such as an object or a server. IceGrid prevents other clients from using the resource until the client releases it or the session expires. Sessions enhance security through the use of an authentication mechanism that can be integrated with a Glacier2 router.

- Automatic failover
  Ice supports automatic retry and failover in any proxy that contains multiple endpoints. When combined with IceGrid's support for replication and load balancing, automatic failover means that a failed request results in a client transparently retrying the request on the next endpoint with the lowest load.

- Dynamic queries
  In addition to transparent binding, applications can interact directly with IceGrid to locate objects in a variety of ways.

- Status monitoring
  IceGrid supports Slice interfaces that allow applications to monitor its activities and receive notifications about significant events, enabling the development of custom tools or the integration of IceGrid status events into an existing management framework.

- Administration
  IceGrid includes command-line and graphical administration tools. They are available on all supported platforms and allow you to start, stop, monitor, and reconfigure any server managed by IceGrid.

- Deployment
  Using XML files, you can describe the servers to be deployed on each computer. Templates simplify the description of identical servers.

As grid computing enters the mainstream and compute servers become commodities, users expect more value from their applications. IceGrid, in cooperation with the Ice run time, relieves you of these low-level tasks to accelerate the construction and simplify the administration of your distributed applications.

Copyright 2016, ZeroC, Inc.

Topics

- IceGrid Architecture
- Getting Started with IceGrid
- Using IceGrid Deployment
- Well-Known Objects
- IceGrid Templates
- IceBox Integration with IceGrid
- Object Adapter Replication
- Load Balancing
- Resource Allocation using IceGrid Sessions
- Registry Replication
- Application Distribution
- IceGrid Administrative Sessions
- Glacier2 Integration with IceGrid
- IceGrid XML Reference
- Using Descriptor Variables and Parameters
- IceGrid Property Set Semantics
- IceGrid XML Features
- IceGrid Server Reference
- IceGrid and the Administrative Facility
- Securing IceGrid
- icegridadmin Command Line Tool
- IceGrid Admin Graphical Tool
- IceGrid Server Activation
- IceGrid Troubleshooting
- IceGrid Database Utility

Copyright 2016, ZeroC, Inc.

IceGrid Architecture

An IceGrid domain consists of a *registry* and any number of *nodes*. Together, the registry and nodes cooperate to manage the information and server processes that comprise *applications*. Each application assigns servers to particular nodes. The registry maintains a persistent record of this information, while the nodes are responsible for starting and monitoring their assigned server processes. In a typical configuration, one node runs on each computer that hosts Ice servers. The registry does not consume much processor time, so it commonly runs on the same computer as a node; in fact, the registry and a node can run in the same process if desired. If fault tolerance is desired, the registry supports replication using a master-slave design.

On this page:

- Architecture of a Simple IceGrid Application
- Server Replication with IceGrid
- Deploying an IceGrid Application

Architecture of a Simple IceGrid Application

As an example, this illustration shows a very simple IceGrid application running on a network of three computers. The IceGrid registry is the only process of interest on host PC1, while IceGrid nodes are running on the hosts PC2 and PC3. In this sample application, one server has been assigned to each node.

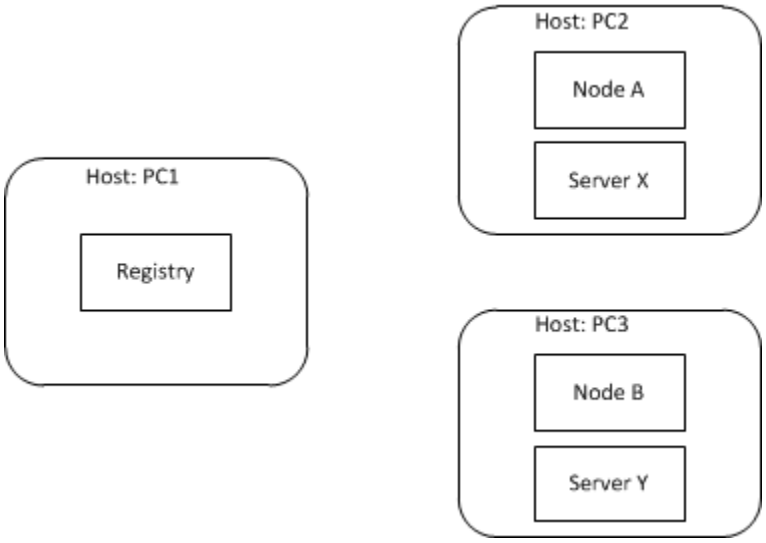

Simple IceGrid application.

From a client application's perspective, the primary responsibility of the registry is to resolve indirect proxies as an Ice location service. As such, this contribution is largely transparent: when a client first attempts to use an indirect proxy, the Ice run time in the client contacts the registry to convert the proxy's symbolic information into endpoints that allow the client to establish a connection.

Although the registry might sound like nothing more than a simple lookup table, reality is quite different. For example, behind the scenes, a locate request might prompt a node to start the target server automatically, or the registry might select appropriate endpoints based on load statistics from each computer.

This also illustrates the benefits of indirect proxies: the location service can provide a great deal of functionality without any special action by the client and, unlike with direct proxies, the client does not need advance knowledge of the address and port of a server. The extra level of indirection adds some latency to the client's first use of a proxy; however, all subsequent interactions occur directly between client and server, so the cost is negligible. Furthermore, indirection allows servers to migrate to different computers without the need to update proxies held by clients.

Server Replication with IceGrid

IceGrid's flexibility allows an endless variety of configurations. For example, suppose we have a grid network and want to replicate a server on each blade, as shown below:

Copyright 2016, ZeroC, Inc.

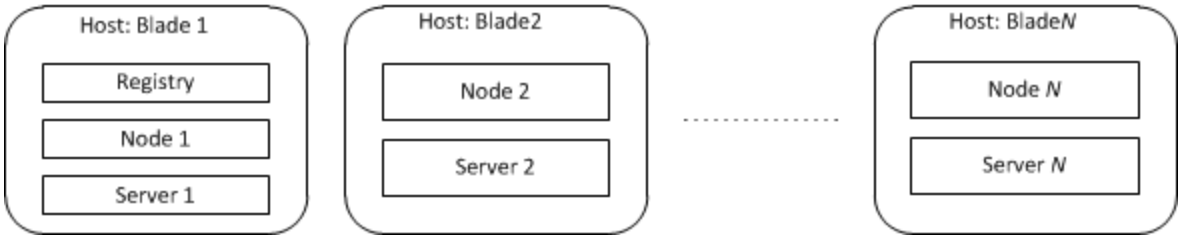

Replicated server on grid network.

Replication in Ice is based on object adapters, not servers. Any object adapter in any server could participate in replication, but it is far more likely that all of the replicated object adapters are created by instances of the same server executable that is running on each computer. We are using this configuration in the example shown above, but IceGrid requires each server to have a unique name. `Server 1` and `Server 2` are our unique names for the same executable.

The binding process works somewhat differently when replication is involved, since the registry now has multiple object adapters to choose from. The description of the IceGrid application drives the registry's decision about which object adapter (or object adapters) to use. For example, the registry could consider the system load of each computer (as periodically reported by the nodes) and return the endpoints of the object adapter on the computer with the lowest load. It is also possible for the registry to combine the endpoints of several object adapters, in which case the Ice run time in the client would select the endpoint for the initial connection attempt.

Deploying an IceGrid Application

In IceGrid, *deployment* is the process of describing an application to the registry. This description includes the following information:

- Replica groups
  A *replica group* is the term for a collection of replicated object adapters. An application can create any number of replica groups. Each group requires a unique identifier.
- Nodes
  An application must assign its servers to one or more nodes.
- Servers
  A server's description includes a unique name and the path to its executable. It also lists the object adapters it creates.
- Object adapters
  Information about an object adapter includes its endpoints and any well-known objects it advertises. If the object adapter is a member of a replica group, it must also supply that group's identifier.
- Objects
  A *well-known object* is one that is known solely by its identity. The registry maintains a global list of such objects for use during locate requests.

IceGrid uses the term *descriptor* to refer to the description of an application and its components; deploying an application involves creating its descriptors in the registry. The are several ways to accomplish this:

- You can use a command-line tool that reads a file containing an XML representation of the descriptors.
- You can create descriptors interactively with the graphical administration tool.
- You can create descriptors programmatically via IceGrid's administrative interface.

The registry server must be running in order to deploy an application, but it is not necessary for nodes to be active. Nodes that are started after deployment automatically retrieve the information they need from the registry. Once deployed, you can update the application at any time.

See Also

- Locators
- Connection Establishment
- Object Adapters
- Object Adapter Replication
- Well-Known Objects

Copyright 2016, ZeroC, Inc.

Getting Started with IceGrid

This page introduces a sample application that will help us demonstrate IceGrid's capabilities.

On this page:

- The Ripper Application
- Initial Ripper Architecture
- Ripper Registry Configuration
- Ripper Client Configuration
- Ripper Server Configuration
- Starting the Registry for the Ripper Application
- Starting the Ripper Server
- Ripper Progress Review

The Ripper Application

Our application "rips" music tracks from a compact disc (CD) and encodes them as MP3 files, as shown below:

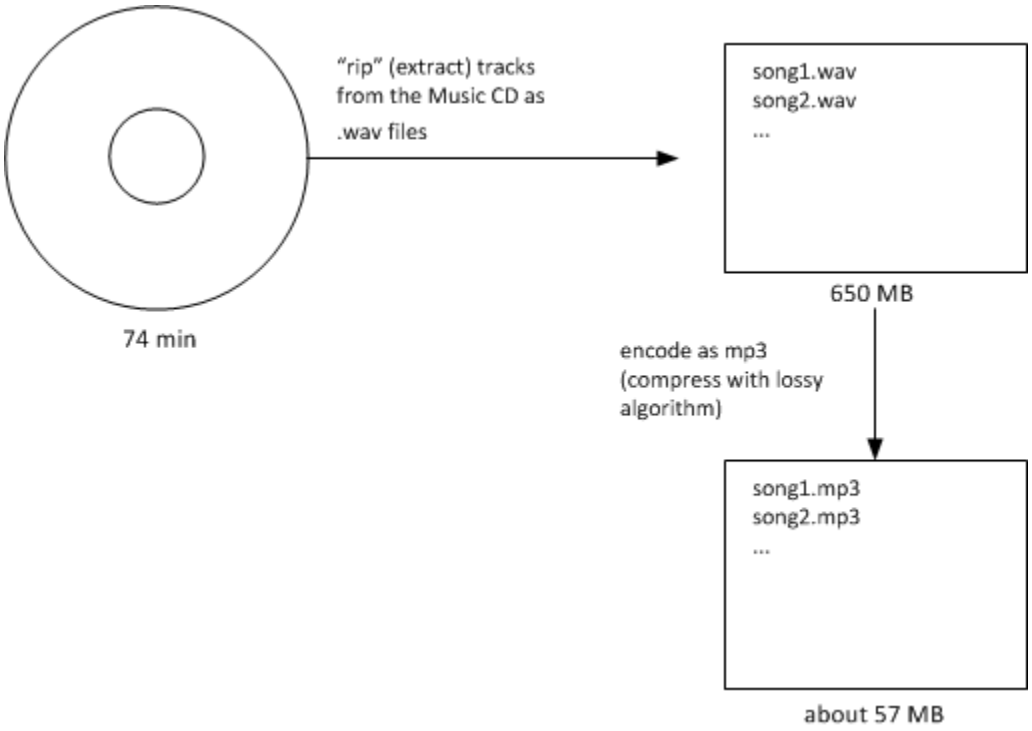

Overview of sample application.

Ripping an entire CD usually takes several minutes because the MP3 encoding requires lots of CPU cycles. Our distributed ripper application accelerates this process by taking advantage of powerful CPUs on remote Ice servers, enabling us to process many songs in parallel.

The Slice interface for the MP3 encoder is straightforward:

Copyright 2016, ZeroC, Inc.

Slice

```
#include <Ice/BuiltinSequences.ice>
module Ripper {
exception EncodingFailedException {
    string reason;
};

sequence<short> Samples;

interface Mp3Encoder {
    // Input: PCM samples for left and right channels
    // Output: MP3 frame(s).
    Ice::ByteSeq encode(Samples leftSamples, Samples rightSamples)
        throws EncodingFailedException;

    // You must flush to get the last frame(s). Flush also
    // destroys the encoder object.
    Ice::ByteSeq flush()
        throws EncodingFailedException;
};

interface Mp3EncoderFactory
{
    Mp3Encoder* createEncoder();
};
};
```

The implementation of the encoding algorithm is not relevant for the purposes of this discussion. Instead, we will focus on incrementally improving the application as we discuss IceGrid features.

Initial Ripper Architecture

The initial architecture for our application is intentionally simple, consisting of an IceGrid registry and a server that we start manually. This illustration shows how the client's invocation on its `EncoderFactory` proxy causes an implicit locate request:

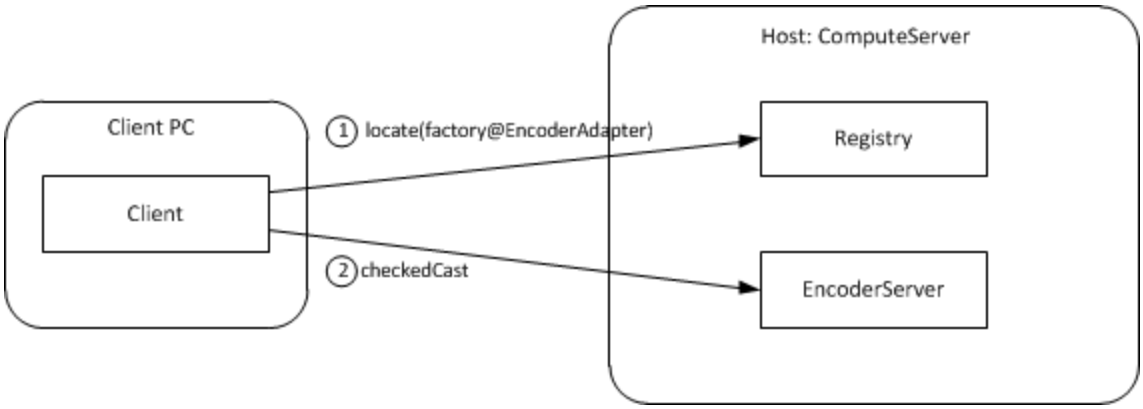

Initial architecture for the ripper application.

The corresponding C++ code for the client is presented below:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx proxy = communicator->stringToProxy("factory@EncoderAdapter");
Ripper::MP3EncoderFactoryPrx factory =
Ripper::MP3EncoderFactoryPrx::checkedCast(proxy);
Ripper::MP3EncoderPrx encoder = factory->createEncoder();
```

Notice that the client uses an indirect proxy for the `MP3EncoderFactory` object. This stringified proxy can be read literally as "the object with identity `factory` in the object adapter identified as `EncoderAdapter`." The encoding server creates this object adapter and ensures that the object adapter uses this identifier. Since each object adapter must be uniquely identified, the registry can easily determine the server that created the adapter and return an appropriate endpoint to the client.

The client's call to `checkedCast` is the first remote invocation on the factory object, and therefore the locate request is performed during the completion of this invocation. The subsequent call to `createEncoder` is sent directly to the server without further involvement by IceGrid.

Ripper Registry Configuration

The registry needs a subdirectory in which to create its databases, and we will use `/opt/ripper/registry` for this purpose (the directory must exist before starting the registry). We also need to create an Ice configuration file to hold properties required by the registry. The file `/opt/ripper/registry.cfg` contains the following properties:

```
IceGrid.Registry.Client.Endpoints=tcp -p 4061
IceGrid.Registry.Server.Endpoints=tcp
IceGrid.Registry.Internal.Endpoints=tcp
IceGrid.Registry.AdminPermissionsVerifier=IceGrid/NullPermissionsVerifier
IceGrid.Registry.Data=/opt/ripper/registry
IceGrid.Registry.DynamicRegistration=1
```

Several of the properties define endpoints, but only the value of `IceGrid.Registry.Client.Endpoints` needs a fixed port. This property specifies the endpoints of the IceGrid locator service; IceGrid clients must include these endpoints in their definition of `Ice.Default.Locator`, as discussed in the next section. The TCP port number (4061) used in this example has been reserved by the Internet Assigned Numbers Authority (IANA) for the IceGrid registry, along with SSL port number 4062.

Several other properties are worth mentioning:

- `IceGrid.Registry.AdminPermissionsVerifier`
  Controls access to the registry's administrative functionality.
- `IceGrid.Registry.Data`
  Specifies the registry's database directory.
- `IceGrid.Registry.DynamicRegistration`
  If set to a non-zero value, allows servers to register their object adapters. Dynamic registration is explained in more detail below.

By default, IceGrid will not permit a server to register its object adapters without using IceGrid's deployment facility. In some situations, such as in this sample application, you may want a client to be able to bind indirectly to a server without having to first deploy the server. That is, simply starting the server should be sufficient to make the server register itself with IceGrid and be reachable from clients.

You can achieve this by running the registry with the property `IceGrid.Registry.DynamicRegistration` set to a non-zero value. With this setting, IceGrid permits an adapter to register itself upon activation even if it has not been previously deployed. To force the server to register its adapters, you must define `Ice.Default.Locator` (so the server can find the registry) and, for each adapter that you wish to register, you must set `<adapter-name>.AdapterId` to an identifier that is unique within the registry. Setting the `<adapter-name>.AdapterId` property also causes the adapter to no longer create direct proxies but rather to create indirect proxies that clients must resolve via the registry.

Copyright 2016, ZeroC, Inc.

Ripper Client Configuration

The client requires only minimal configuration, namely a value for the property `Ice.Default.Locator`. This property supplies the Ice run time with the proxy for the locator service. In IceGrid, the locator service is implemented by the registry, and the locator object is available on the registry's client endpoints. The property `IceGrid.Registry.Client.Endpoints` defined above provides most of the information we need to construct the proxy. The missing piece is the identity of the locator object, which defaults to `IceGrid/Locator` but may change based on the registry's configuration:

```
Ice.Default.Locator=IceGrid/Locator:tcp -h registryhost -p 4061
```

The use of a locator service allows the client to take advantage of indirect binding and avoid static dependencies on server endpoints. However, the locator proxy must have a fixed port, otherwise the client has a bootstrapping problem: it cannot resolve indirect proxies without knowing the endpoints of the locator service.

IceLocatorDiscovery eliminates the need for a client to define the `Ice.Default.Locator` property by using UDP multicast to discover registries at run time.

Ripper Server Configuration

We use `/opt/ripper/server.cfg` as the server's configuration file. It contains the following properties:

```
EncoderAdapter.AdapterId=EncoderAdapter
EncoderAdapter.Endpoints=tcp
Ice.Default.Locator=IceGrid/Locator:tcp -h registryhost -p 4061
```

The properties are described below:

- `EncoderAdapter.AdapterId`
  This property supplies the object adapter identifier that the client uses in its indirect proxy (e.g., `factory@EncoderAdapter`).
- `EncoderAdapter.Endpoints`
  This property defines the object adapter's endpoint. Notice that the value does not contain any port information, meaning that the adapter uses a system-assigned port. Without IceGrid, the use of a system-assigned port would pose a significant problem: how would a client create a direct proxy if the adapter's port could change every time the server is restarted? IceGrid solves this problem nicely because clients can use indirect proxies that contain no endpoint dependencies. The registry resolves indirect proxies using the endpoint information supplied by object adapters each time they are activated.
- `Ice.Default.Locator`
  The server requires a value for this property in order to register its object adapter.

Starting the Registry for the Ripper Application

Now that the configuration file is written and the directory structure is prepared, we are ready to start the IceGrid registry:

```
$ icegridregistry --Ice.Config=/opt/ripper/registry.cfg
```

Additional command line options are supported, including those that allow the registry to run as a Windows service or Unix daemon.

Starting the Ripper Server

With the registry up and running, we can now start the server. At a command prompt, we run the program and pass an `--Ice.Config` option indicating the location of the configuration file:

Copyright 2016, ZeroC, Inc.

```
$ /opt/ripper/bin/server --Ice.Config=/opt/ripper/server.cfg
```

Ripper Progress Review

This example demonstrated how to use IceGrid's location service, which is a core component of IceGrid's feature set. By incorporating IceGrid into our application, the client is now able to locate the `MP3EncoderFactory` object using only an indirect proxy and a value for `Ice.Default.Locator`. Furthermore, we can reconfigure the application in any number of ways without modifying the client's code or configuration.

For some applications, the functionality we have already achieved using IceGrid may be entirely sufficient. However, we have only just begun to explore IceGrid's capabilities, and there is much we can still do to improve our application. The next section shows how we can avoid the need to start our server manually by deploying our application onto an IceGrid node.

See Also

- Locator Configuration for a Client
- Resource Allocation using IceGrid Sessions
- Well-Known Registry Objects
- Using IceGrid Deployment
- Object Adapter Endpoints
- icegridregistry
- IceGrid.*

Copyright 2016, ZeroC, Inc.

Using IceGrid Deployment

Here we extend the capabilities of our sample application using IceGrid's deployment facility.

On this page:

- Ripper Architecture using Deployment
- Ripper Deployment Descriptors
- Ripper Registry and Node Configuration
- Ripper Server Configuration using Deployment
- Starting the Node for the Ripper Application
- Deploying the Ripper Application
- Ripper Progress Review
- Adding Nodes to the Ripper Application
  - Descriptor Changes
  - Configuration Changes
  - Redeploying the Application
  - Client Changes

Ripper Architecture using Deployment

The revised architecture for our application consists of a single IceGrid node responsible for our encoding server that runs on the computer named `ComputeServer`. The illustration below shows the client's initial invocation on its indirect proxy and the actions that IceGrid takes to make this invocation possible:

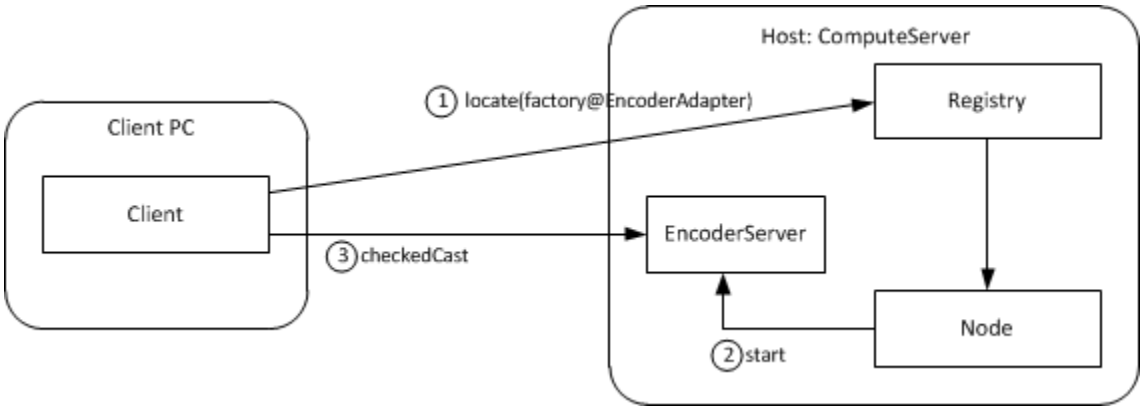

Architecture for deployed ripper application.

In contrast to the initial architecture, we no longer need to manually start our server. In this revised application, the client's locate request prompts the registry to query the node about the server's state and start it if necessary. Once the server starts successfully, the locate request completes and subsequent client communication occurs directly with the server.

Ripper Deployment Descriptors

We can deploy our application using the `icegridadmin` command line utility, but first we must define our descriptors in XML. The descriptors are quite brief:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="Ripper">
        <node name="Node1">
            <server id="EncoderServer" exe="/opt/ripper/bin/server"
activation="on-demand">
                <adapter name="EncoderAdapter" id="EncoderAdapter"
endpoints="tcp"/>
            </server>
        </node>
    </application>
</icegrid>
```

For IceGrid's purposes, we have named our application `Ripper`. It consists of a single server, `EncoderServer`, assigned to the node `Node1`.

Since a computer typically runs only one node process, you might be tempted to give the node a name that identifies its host (such as `ComputeServerNode`). However, this naming convention becomes problematic as soon as you need to migrate the node to another host.

The server's `exe` attribute supplies the pathname of its executable, and the `activation` attribute indicates that the server should be activated on demand when necessary.

The object adapter's descriptor is the most interesting. As you can see, the `name` and `id` attributes both specify the value `EncoderAdapter`. The value of `name` reflects the adapter's name in the server process (i.e., the argument passed to `createObjectAdapter`) that is used for configuration purposes, whereas the value of `id` uniquely identifies the adapter within the registry and is used in indirect proxies. These attributes are not required to have the same value. Had we omitted the `id` attribute, IceGrid would have composed a unique value by combining the server name and adapter name to produce the following identifier:

```
EncoderServer.EncoderAdapter
```

The `endpoints` attribute defines one or more endpoints for the adapter. As explained earlier, these endpoints do not require a fixed port.

Refer to the XML reference for detailed information on using XML to define descriptors.

Ripper Registry and Node Configuration

In our initial registry configuration, we created the directory `/opt/ripper/registry` for use by the registry. The node also needs a subdirectory for its own purposes, so we will use `/opt/ripper/node`. Again, these directories must exist before starting the registry and node.

We also need to create an Ice configuration file to hold properties required by the registry and node. The file `/opt/ripper/config` contains the following properties:

Copyright 2016, ZeroC, Inc.

```
# Registry properties
IceGrid.Registry.Client.Endpoints=tcp -p 4061
IceGrid.Registry.Server.Endpoints=tcp
IceGrid.Registry.Internal.Endpoints=tcp
IceGrid.Registry.AdminPermissionsVerifier=IceGrid/NullPermissionsVerifier
IceGrid.Registry.Data=/opt/ripper/registry

# Node properties
IceGrid.Node.Endpoints=tcp
IceGrid.Node.Name=Node1
IceGrid.Node.Data=/opt/ripper/node
IceGrid.Node.CollocateRegistry=1
Ice.Default.Locator=IceGrid/Locator:tcp -p 4061
```

The registry and node can share this configuration file. In fact, by enabling `IceGrid.Node.CollocateRegistry`, we have indicated that the registry and node should run in the same process.

One difference from our initial configuration is that we no longer define `IceGrid.Registry.DynamicRegistration`. By omitting this property, we force the registry to reject the registration of object adapters that have not been deployed.

The node properties are explained below:

- `IceGrid.Node.Endpoints`
  This property specifies the node's endpoints. A fixed port is not required.
- `IceGrid.Node.Name`
  This property defines the unique name for this node. Its value must match the descriptor we wrote above.
- `IceGrid.Node.Data`
  This property specifies the node's data directory.
- `Ice.Default.Locator`
  This property is defined for use by the `icegridadmin` tool. The node would also require this property if the registry is not collocated. Refer to our discussion of the ripper client configuration for more information on this setting.

Ripper Server Configuration using Deployment

Server configuration is accomplished using descriptors. During deployment, the node creates a subdirectory tree for each server. Inside this tree the node creates a configuration file containing properties derived from the server's descriptors. For instance, the adapter's descriptor ge nerates the following properties in the server's configuration file:

```
# Server configuration
Ice.Admin.ServerId=EncoderServer
Ice.Admin.Endpoints=tcp -h 127.0.0.1
Ice.ProgramName=EncoderServer
# Object adapter EncoderAdapter
EncoderAdapter.Endpoints=tcp
EncoderAdapter.AdapterId=EncoderAdapter
Ice.Default.Locator=IceGrid/Locator:default -p 4061
```

As you can see, the configuration file that IceGrid generates from the descriptor resembles the initial configuration, with two additional properties:

- `Ice.Admin.ServerId`
- `Ice.Admin.Endpoints`

Copyright 2016, ZeroC, Inc.

The `Ice.Admin.Endpoints` property enables the administrative facility that, among other features, allows an IceGrid node to gracefully deactivate the server.

Using the directory structure we established for our ripper application, the configuration file for `EncoderServer` has the file name shown below:

```
/opt/ripper/node/servers/EncoderServer/config/config
```

Note that this file should not be edited directly because any changes you make are lost the next time the node regenerates the file. The correct way to add properties to the file is to include property definitions in the server's descriptor. For example, we can add the property `Ice.Trace.Network=1` by modifying the server descriptor as follows:

```
<icegrid>
    <application name="Ripper">
        <node name="Node1">
            <server id="EncoderServer" exe="/opt/ripper/bin/server"
activation="on-demand">
                <adapter name="EncoderAdapter" id="EncoderAdapter"
endpoints="tcp"/>
                <property name="Ice.Trace.Network" value="1"/>
            </server>
        </node>
    </application>
</icegrid>
```

When a node activates a server, it passes the location of the server's configuration file using the `--Ice.Config` command-line argument. If you start a server manually from a command prompt, you must supply this argument yourself.

Starting the Node for the Ripper Application

Now that the configuration file is written and the directory structure is prepared, we are ready to start the IceGrid registry and node. Using a collocated registry and node, we only need to use one command:

```
$ icegridnode --Ice.Config=/opt/ripper/config
```

Additional command line options are supported, including those that allow the node to run as a Windows service or Unix daemon.

Deploying the Ripper Application

With the registry up and running, it is now time to deploy our application. Like our client, the `icegridadmin` utility also requires a definition for the `Ice.Default.Locator` property. We can start the utility with the following command:

```
$ icegridadmin --Ice.Config=/opt/ripper/config
```

After confirming that it can contact the registry, `icegridadmin` provides a command prompt at which we deploy our application. Assuming our descriptor is stored in `/opt/ripper/app.xml`, the deployment command is shown below:

Copyright 2016, ZeroC, Inc.

```
>>> application add "/opt/ripper/app.xml"
```

Next, confirm that the application has been deployed:

```
>>> application list
Ripper
```

You can start the server using this command:

```
>>> server start EncoderServer
```

Finally, you can retrieve the current endpoints of the object adapter:

```
>>> adapter endpoints EncoderAdapter
```

If you want to experiment further using `icegridadmin`, issue the `help` command and review the available commands.

Ripper Progress Review

We have deployed our first IceGrid application, but you might be questioning whether it was worth the effort. Even at this early stage, we have already gained several benefits:

- We no longer need to manually start the encoder server before starting the client, because the IceGrid node automatically starts it if it is not active at the time a client needs it. If the server happens to terminate for any reason, such as an IceGrid administrative action or a server programming error, the node restarts it without intervention on our part.
- We can manage the application remotely using one of the IceGrid administration tools. The ability to remotely modify applications, start and stop servers, and inspect every aspect of your configuration is a significant advantage.

Admittedly, we have not made much progress yet in our stated goal of improving the performance of the ripper over alternative solutions that are restricted to running on a single computer. Our client now has the ability to easily delegate the encoding task to a server running on another computer, but we have not achieved the parallelism that we really need. For example, if the client created a number of encoders and used them simultaneously from multiple threads, the encoding performance might actually be *worse* than simply encoding the data directly in the client, as the remote computer would likely slow to a crawl while attempting to task-switch among a number of processor-intensive tasks.

Adding Nodes to the Ripper Application

Adding more nodes to our environment would allow us to distribute the encoding load to more compute servers. Using the techniques we have learned so far, let us investigate the impact that adding a node would have on our descriptors, configuration, and client application.

Descriptor Changes

The addition of a node is mainly an exercise in cut and paste:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="Ripper">
        <node name="Node1">
            <server id="EncoderServer1" exe="/opt/ripper/bin/server"
activation="on-demand">
                <adapter name="EncoderAdapter" endpoints="tcp"/>
            </server>
        </node>
        <node name="Node2">
            <server id="EncoderServer2" exe="/opt/ripper/bin/server"
activation="on-demand">
                <adapter name="EncoderAdapter" endpoints="tcp"/>
            </server>
        </node>
    </application>
</icegrid>
```

Note that we now have two `node` elements instead of a single one. You might be tempted to simply use the host name as the node name. However, in general, that is not a good idea. For example, you may want to run several IceGrid nodes on a single machine (for example, for testing). Similarly, you may have to rename a host at some point, or need to migrate a node to a different host. But, unless you also rename the node, that leads to the situation where you have a node with the name of a (possibly obsolete) host when the node in fact is not running on that host. Obviously, this makes for a confusing configuration — it is better to use abstract node names, such as `Node1`.

Aside from the new `node` element, notice that the server identifiers must be unique. The adapter name, however, can remain as `EncoderAdapter` because this name is used only for local purposes within the server process. In fact, using a different name for each adapter would actually complicate the server implementation, since it would somehow need to discover the name it should use when creating the adapter.

We have also removed the `id` attribute from our adapter descriptors; the default values supplied by IceGrid are sufficient for our purposes.

Configuration Changes

We can continue to use the configuration file we created earlier for our combined registry-node process. We need a separate configuration file for `Node2`, primarily to define a different value for the property `IceGrid.Node.Name`. However, we also cannot have two nodes configured with `IceGrid.Node.CollocateRegistry` because only one master registry is allowed, so we must remove this property:

```
IceGrid.Node.Endpoints=tcp
IceGrid.Node.Name=Node2
IceGrid.Node.Data=/opt/ripper/node

Ice.Default.Locator=IceGrid/Locator:tcp -h registryhost -p 4061
```

We assume that `/opt/ripper/node` refers to a local file system directory on the computer hosting `Node2`, and not a shared volume, because two nodes must not share the same data directory.

We have also modified the locator proxy to include the address of the host on which the registry is running.

Using IceLocatorDiscovery allows a node to discover its registry at run time without the need to define `Ice.Default.Locator`.

Redeploying the Application

Copyright 2016, ZeroC, Inc.

After saving the new descriptors, you need to redeploy the application. Using `icegridadmin`, issue the following command:

```
$ icegridadmin --Ice.Config=/opt/ripper/config
>>> application update "/opt/ripper/app.xml"
```

If an update affects any of the application's servers that are currently running, IceGrid automatically stops those servers prior to performing the update and restarts them again after the update is complete. We can determine whether an update would require any restarts using the `application diff` command:

```
$ icegridadmin --Ice.Config=/opt/ripper/config
>>> application diff --servers "/opt/ripper/app.xml"
```

To ensure that our update does not impact any active servers, we can use the `--no-restart` option:

```
$ icegridadmin --Ice.Config=/opt/ripper/config
>>> application update --no-restart "/opt/ripper/app.xml"
```

With this option, the update would fail if any servers required a restart.

Client Changes

We have added a new node, but we still need to modify our client to take advantage of it. As it stands now, our client can delegate an encoding task to one of the two `MP3EncoderFactory` objects. The client selects a factory by using the appropriate indirect proxy:

- `factory@EncoderServer1.EncoderAdapter`
- `factory@EncoderServer2.EncoderAdapter`

In order to distribute the tasks among both factories, the client could use a random number generator to decide which factory receives the next task:

C++

```
string adapter;
if ((rand() % 2) == 0)
    adapter = "EncoderServer1.EncoderAdapter";
else
    adapter = "EncoderServer2.EncoderAdapter";
Ice::ObjectPrx proxy = communicator->stringToProxy("factory@" + adapter);
Ripper::MP3EncoderFactoryPrx factory =
Ripper::MP3EncoderFactoryPrx::checkedCast(proxy);
Ripper::MP3EncoderPrx encoder = factory->createEncoder();
```

There are a few disadvantages in this design:

- The client application must be modified each time a new compute server is added or removed because it knows all of the adapter identifiers.
- The client cannot distribute the load intelligently; it is just as likely to assign a task to a heavily-loaded computer as it is an idle one.

We describe better solutions in the sections that follow.

Copyright 2016, ZeroC, Inc.

See Also

- IceGrid Server Activation
- Creating an Object Adapter
- Object Adapter Endpoints
- Getting Started with IceGrid
- icegridadmin Command Line Tool
- IceGrid and the Administrative Facility
- icegridnode
- Adapter Descriptor Element
- IceGrid.*

Copyright 2016, ZeroC, Inc.

Well-Known Objects

On this page:

- Overview of Well-Known Objects
- Well-Known Object Types
- Deploying Well-Known Objects
- Adding Well-Known Objects Programmatically
- Adding Well-Known Objects with icegridadmin
- Querying Well-Known Objects
- Using Well-Known Objects in the Ripper Application
  - Adding Well-Known Objects to the Ripper Deployment
  - Querying Ripper Objects with findAllObjectsByType
  - Querying Ripper Objects with findObjectByType
  - Querying Ripper Objects with findObjectByTypeOnLeastLoadedNode
  - Ripper Progress Review

Overview of Well-Known Objects

There are two types of indirect proxies: one specifies an identity and an object adapter identifier, while the other contains only an identity. The latter type of indirect proxy is known as a *well-known proxy*. A well-known proxy refers to a well-known object, that is, its identity alone is sufficient to allow the client to locate it. Ice requires all object identities in an application to be unique, but typically only a select few objects are able to be located only by their identities.

In earlier sections we showed the relationship between indirect proxies containing an object adapter identifier and the IceGrid configuration. Briefly, in order for a client to use a proxy such as `factory@EncoderAdapter`, an object adapter must be given the identifier `EncoderAdapter`.

A similar requirement exists for well-known objects. The registry maintains a table of these objects, which can be populated in a number of ways:

- statically in descriptors,
- programmatically using IceGrid's administrative interface,
- dynamically using an IceGrid administration tool.

The registry's database maps an object identity to a proxy. A locate request containing only an identity prompts the registry to consult this database. If a match is found, the registry examines the associated proxy to determine if additional work is necessary. For example, consider the well-known objects in the following table.

Identity	Proxy
`Object1`	`Object1:tcp -p 10001`
`Object2`	`Object2@TheAdapter`
`Object3`	`Object3`

The proxy associated with `Object1` already contains endpoints, so the registry can simply return this proxy to the client.

For `Object2`, the registry notices the adapter ID and checks to see whether it knows about an adapter identified as `TheAdapter`. If it does, it attempts to obtain the endpoints of that adapter, which may cause its server to be started. If the registry is successfully able to determine the adapter's endpoints, it returns a direct proxy containing those endpoints to the client. If the registry does not recognize `TheAdapter` or cannot obtain its endpoints, it returns the indirect proxy `Object2@TheAdapter` to the client. Upon receipt of another indirect proxy, the Ice run time in the client will try once more to resolve the proxy, but generally this will not succeed and the Ice run time in the client will raise a `NoEndpointException` as a result.

Finally, `Object3` represents a hopeless situation: how can the registry resolve `Object3` when its associated proxy refers to itself? In this case, the registry returns the proxy `Object3` to the client, which causes the client to raise `NoEndpointException`. Clearly, you should avoid this situation.

Well-Known Object Types

The registry's database not only associates an identity with a proxy, but also a type. Technically, the "type" is an arbitrary string but, by convention, that string represents the most-derived Slice type of the object. For example, the Slice type ID of the encoder factory in our ripper application is `::Ripper::MP3EncoderFactory`.

Object types are useful when performing queries.

Copyright 2016, ZeroC, Inc.

Deploying Well-Known Objects

The `object` descriptor adds a well-known object to the registry. It must appear within the context of an adapter descriptor, as shown in the XML example below:

XML

```
<icegrid>
    <application name="Ripper">
        <node name="Node1">
            <server id="EncoderServer" exe="/opt/ripper/bin/server"
activation="on-demand">
                <adapter name="EncoderAdapter" id="EncoderAdapter" endpoints="tcp">
                    <object identity="EncoderFactory"
type="::Ripper::MP3EncoderFactory"/>
                </adapter>
            </server>
        </node>
    </application>
</icegrid>
```

During deployment, the registry associates the identity `EncoderFactory` with the indirect proxy `EncoderFactory@EncoderAdapter`. If the adapter descriptor had omitted the adapter ID, the registry would have generated a unique identifier by combining the server ID and the adapter name.

In this example, the object's type is specified explicitly.

Adding Well-Known Objects Programmatically

The `IceGrid::Admin` interface defines several operations that manipulate the registry's database of well-known objects:

Slice

```
module IceGrid {
interface Admin {
    ...
    void addObject(Object* obj)
        throws ObjectExistsException,
               DeploymentException;
    void updateObject(Object* obj)
        throws ObjectNotRegisteredException,
               DeploymentException;
    void addObjectWithType(Object* obj, string type)
        throws ObjectExistsException,
               DeploymentException;
    void removeObject(Ice::Identity id)
        throws ObjectNotRegisteredException,
               DeploymentException;
    ...
};
};
```


Copyright 2016, ZeroC, Inc.

- `addObject`
  The `addObject` operation adds a new object to the database. The proxy argument supplies the identity of the well-known object. If an object with the same identity has already been registered, the operation raises `ObjectExistsException`. Since this operation does not accept an argument supplying the object's type, the registry invokes `ice_id` on the given proxy to determine its most-derived type. The implication here is that the object must be available in order for the registry to obtain its type. If the object is not available, `addObject` raises `DeploymentException`.
- `updateObject`
  The `updateObject` operation supplies a new proxy for the well-known object whose identity is encapsulated by the proxy. If no object with the given identity is registered, the operation raises `ObjectNotRegisteredException`. The object's type is not modified by this operation.
- `addObjectWithType`
  The `addObjectWithType` operation behaves like `addObject`, except the object's type is specified explicitly and therefore the registry does not attempt to invoke `ice_id` on the given proxy (even if the type is an empty string).
- `removeObject`
  The `removeObject` operation removes the well-known object with the given identity from the database. If no object with the given identity is registered, the operation raises `ObjectNotRegisteredException`.

The following C++ example produces the same result as the descriptor we deployed earlier:

C++

```
Ice::ObjectAdapterPtr adapter =
communicator->createObjectAdapter("EncoderAdapter");
Ice::Identity ident = communicator->stringToIdentity("EncoderFactory");
FactoryPtr f= new FactoryI;
Ice::ObjectPrx factory = adapter->add(f, ident);
IceGrid::AdminPrx admin = // ...
try {
    admin->addObject(factory); // OOPS!
} catch (const IceGrid::ObjectExistsException &) {
    admin->updateObject(factory);
}
```

After obtaining a proxy for the `IceGrid::Admin` interface, the code invokes `addObject`. Notice that the code traps `ObjectExistsException` and calls `updateObject` instead when the object is already registered.

There is one subtle problem in this code: calling `addObject` causes the registry to invoke `ice_id` on our factory object, but we have not yet activated the object adapter. As a result, our program will hang indefinitely at the call to `addObject`. One solution is to activate the adapter prior to the invocation of `addObject`; another solution is to use `addObjectWithType` as shown below:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectAdapterPtr adapter =
communicator->createObjectAdapter("EncoderAdapter");
Ice::Identity ident = communicator->stringToIdentity("EncoderFactory");
FactoryPtr f = new FactoryI;
Ice::ObjectPrx factory = adapter->add(f, ident);
IceGrid::AdminPrx admin = // ...
try {
    admin->addObjectWithType(factory, factory->ice_id());
} catch (const IceGrid::ObjectExistsException &) {
    admin->updateObject(factory);
}
```

Adding Well-Known Objects with `icegridadmin`

The icegridadmin utility provides commands that are the functional equivalents of the Slice operations for managing well-known objects. We can use the utility to manually register the `EncoderFactory` object from our descriptors:

```
$ icegridadmin --Ice.Config=/opt/ripper/config
>>> object add "EncoderFactory@EncoderAdapter"
```

Use the `object list` command to verify that the object was registered successfully:

```
>>> object list
EncoderFactory
IceGrid/Query
IceGrid/Locator
IceGrid/Registry
IceGrid/InternalRegistry-Master
```

To specify the object's type explicitly, append it to the `object add` command:

```
>>> object add "EncoderFactory@EncoderAdapter" "::Ripper::MP3EncoderFactory"
```

Finally, the object is removed from the registry like this:

```
>>> object remove "EncoderFactory"
```

Querying Well-Known Objects

Copyright 2016, ZeroC, Inc.

The registry's database of well-known objects is not used solely for resolving indirect proxies. The database can also be queried interactively to find objects in a variety of ways. The `IceGrid::Query` interface supplies this functionality:

Slice

```
module IceGrid {
enum LoadSample {
    LoadSample1,
    LoadSample5,
    LoadSample15
};

interface Query {
    idempotent Object* findObjectById(Ice::Identity id);
    idempotent Object* findObjectByType(string type);
    idempotent Object* findObjectByTypeOnLeastLoadedNode(string type, LoadSample sa
mple);
    idempotent Ice::ObjectProxySeq findAllObjectsByType(string type);
    idempotent Ice::ObjectProxySeq findAllReplicas(Object* proxy);
};
};
```

- `findObjectById`
  The `findObjectById` operation returns the proxy associated with the given identity of a well-known object. It returns a null proxy if no match was found.
- `findObjectByType`
  The `findObjectByType` operation returns a proxy for an object registered with the given type. If more than one object has the same type, the registry selects one at random. The operation returns a null proxy if no match was found.
- `findObjectByTypeOnLeastLoadedNode`
  The `findObjectByTypeOnLeastLoadedNode` operation considers the system load when selecting one of the objects with the given type. If the registry is unable to determine which node hosts an object (for example, because the object was registered with a direct proxy and not an adapter ID), the object is considered to have a load value of `1` for the purposes of this operation. The sample argument determines the interval over which the loads are averaged (one, five, or fifteen minutes). The operation returns a null proxy if no match was found.
- `findAllObjectsByType`
  The `findAllObjectsByType` operation returns a sequence of proxies representing the well-known objects having the given type. The operation returns an empty sequence if no match was found.
- `findAllReplicas`
  Given an indirect proxy for a replicated object, the `findAllReplicas` operation returns a sequence of proxies representing the individual replicas. An application can use this operation when it is necessary to communicate directly with one or more replicas.

Be aware that the operations accepting a `type` parameter are not equivalent to invoking `ice_isA` on each object to determine whether it supports the given type, a technique that would not scale well for a large number of registered objects. Rather, the operations simply compare the given type to the object's registered type or, if the object was registered without a type, to the object's most-derived Slice type as determined by the registry.

Using Well-Known Objects in the Ripper Application

Well-known objects are another IceGrid feature we can incorporate into our ripper application.

Adding Well-Known Objects to the Ripper Deployment

First we'll modify the descriptors to add two well-known objects:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="Ripper">
        <node name="Node1">
            <server id="EncoderServer1" exe="/opt/ripper/bin/server"
activation="on-demand">
                <adapter name="EncoderAdapter" endpoints="tcp">
                    <object identity="EncoderFactory1"
type="::Ripper::MP3EncoderFactory"/>
                </adapter>
            </server>
        </node>
        <node name="Node2">
            <server id="EncoderServer2" exe="/opt/ripper/bin/server"
activation="on-demand">
                <adapter name="EncoderAdapter" endpoints="tcp">
                    <object identity="EncoderFactory2"
type="::Ripper::MP3EncoderFactory"/>
                </adapter>
            </server>
        </node>
    </application>
</icegrid>
```

At first glance, the addition of the well-known objects does not appear to simplify our client very much. Rather than selecting which of the two adapters receives the next task, we now need to select one of the well-known objects.

Querying Ripper Objects with `findAllObjectsByType`

The `IceGrid::Query` interface provides a way to eliminate the client's dependency on object adapter identifiers and object identities. Since our factories are registered with the same type, we can search for all objects of that type:

C++

```
Ice::ObjectPrx proxy = communicator->stringToProxy("IceGrid/Query");
IceGrid::QueryPrx query = IceGrid::QueryPrx::checkedCast(proxy);
Ice::ObjectProxySeq seq;
string type = Ripper::MP3EncoderFactory::ice_staticId();
seq = query->findAllObjectsByType(type);
if (seq.empty()) {
    // no match
}
Ice::ObjectProxySeq::size_type index = ... // random number
Ripper::MP3EncoderFactoryPrx factory =
Ripper::MP3EncoderFactoryPrx::checkedCast(seq[index]);
Ripper::MP3EncoderPrx encoder = factory->createEncoder();
```

This example invokes `findAllObjectsByType` and then randomly selects an element of the sequence.

Copyright 2016, ZeroC, Inc.

Querying Ripper Objects with `findObjectByType`

We can simplify the client further using `findObjectByType` instead, which performs the randomization for us:

C++

```
Ice::ObjectPrx proxy = communicator->stringToProxy("IceGrid/Query");
IceGrid::QueryPrx query = IceGrid::QueryPrx::checkedCast(proxy);
Ice::ObjectPrx obj;
string type = Ripper::MP3EncoderFactory::ice_staticId();
obj = query->findObjectByType(type);
if (!obj) {
    // no match
}
Ripper::MP3EncoderFactoryPrx factory =
Ripper::MP3EncoderFactoryPrx::checkedCast(obj);
Ripper::MP3EncoderPrx encoder = factory->createEncoder();
```

Querying Ripper Objects with `findObjectByTypeOnLeastLoadedNode`

So far the use of `IceGrid::Query` has allowed us to simplify our client, but we have not gained any functionality. If we replace the call to `findObjectByType` with `findObjectByTypeOnLeastLoadedNode`, we can improve the client by distributing the encoding tasks more intelligently. The change to the client's code is trivial:

C++

```
Ice::ObjectPrx proxy = communicator->stringToProxy("IceGrid/Query");
IceGrid::QueryPrx query = IceGrid::QueryPrx::checkedCast(proxy);
Ice::ObjectPrx obj;
string type = Ripper::MP3EncoderFactory::ice_staticId();
obj = query->findObjectByTypeOnLeastLoadedNode(type,
    IceGrid::LoadSample1);
if (!obj) {
    // no match
}
Ripper::MP3EncoderFactoryPrx factory =
Ripper::MP3EncoderFactoryPrx::checkedCast(obj);
Ripper::MP3EncoderPrx encoder = factory->createEncoder();
```

Ripper Progress Review

Incorporating intelligent load distribution is a worthwhile enhancement and is a capability that would be time consuming to implement ourselves. However, our current design uses only well-known objects in order to make queries possible. We do not really need the encoder factory object on each compute server to be individually addressable as a well-known object, a fact that seems clear when we examine the identities we assigned to them: `EncoderFactory1`, `EncoderFactory2`, and so on. IceGrid's replication features give us the tools we need to improve our design.

See Also

- Terminology
- Type IDs
- Object Descriptor Element
- IceGrid Administrative Sessions

Copyright 2016, ZeroC, Inc.

- icegridadmin Command Line Tool
- Object Adapter Replication

Copyright 2016, ZeroC, Inc.

IceGrid Templates

IceGrid templates simplify the task of creating the descriptors for an application. A template is a parameterized descriptor that you can instantiate as often as necessary, and they are descriptors in their own right. Templates are components of an IceGrid application and therefore they are stored in the registry's database. As such, their use is not restricted to XML files; templates can also be created and instantiated interactively using the graphical administration tool.

You can define templates for server and service descriptors. The focus of this section is server templates; we discuss service descriptors and templates in the context of IceBox integration.

On this page:

- Server Templates
- Template Parameters
- Adding Properties to a Server Instance
- Default Templates
- Using Templates with icegridadmin

Server Templates

You may recall from a previous example that the XML description of our sample application defined two nearly identical servers:

XML

```
<icegrid>
    <application name="Ripper">
        <node name="Node1">
            <server id="EncoderServer1" exe="/opt/ripper/bin/server"
activation="on-demand">
                <adapter name="EncoderAdapter" endpoints="tcp"/>
            </server>
        </node>
        <node name="Node2">
            <server id="EncoderServer2" exe="/opt/ripper/bin/server"
activation="on-demand">
                <adapter name="EncoderAdapter" endpoints="tcp"/>
            </server>
        </node>
    </application>
</icegrid>
```

This example is an excellent candidate for a server template. Equivalent definitions that incorporate a template are shown below:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="Ripper">
        <server-template id="EncoderServerTemplate">
            <parameter name="index"/>
            <server id="EncoderServer${index}" exe="/opt/ripper/bin/server"
activation="on-demand">
                <adapter name="EncoderAdapter" endpoints="tcp"/>
            </server>
        </server-template>
        <node name="Node1">
            <server-instance template="EncoderServerTemplate" index="1"/>
        </node>
        <node name="Node2">
            <server-instance template="EncoderServerTemplate" index="2"/>
        </node>
    </application>
</icegrid>
```

We have defined a server template named `EncoderServerTemplate`. Nested within the `server-template` element is a server descriptor that defines an encoder server. The only difference between this `server` element and our previous example is that it is now parameterized: the template parameter `index` is used to form unique identifiers for the server and its adapter. The symbol `${index}` is replaced with the value of the `index` parameter wherever it occurs.

The template is instantiated by a `server-instance` element, which may be used anywhere that a `server` element is used. The server instance descriptor identifies the template to be instantiated, and supplies a value for the `index` parameter.

Although we have not significantly reduced the length of our XML file, we have made it more readable. And more importantly, deploying this server on additional nodes has become much easier.

Template Parameters

Parameters enable you to customize each instance of a template as necessary. The example above defined the `index` parameter with a different value for each instance to ensure that identifiers are unique. A parameter may also declare a default value that is used in the template if no value is specified for it. In our sample application the `index` parameter is considered mandatory and therefore should not have a default value, but we can illustrate this feature in another way. For example, suppose that the path name of the server's executable may change on each node. We can supply a default value for this attribute and override it when necessary:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="Ripper">
        <server-template id="EncoderServerTemplate">
            <parameter name="index"/>
            <parameter name="exepath" default="/opt/ripper/bin/server"/>
            <server id="EncoderServer${index}" exe="${exepath}"
activation="on-demand">
                <adapter name="EncoderAdapter" endpoints="tcp"/>
            </server>
        </server-template>
        <node name="Node1">
            <server-instance template="EncoderServerTemplate" index="1"/>
        </node>
        <node name="Node2">
            <server-instance template="EncoderServerTemplate" index="2"
                exepath="/opt/ripper-test/bin/server"/>
        </node>
    </application>
</icegrid>
```

As you can see, the instance on `Node1` uses the default value for the new parameter `exepath`, but the instance on `Node2` defines a different location for the server's executable.

Understanding the semantics of descriptor variables and parameters will help you add flexibility to your own IceGrid applications.

Adding Properties to a Server Instance

As we saw in the preceding section, template parameters allow you to customize each instance of a server template, and template parameters with default values allow you to define commonly used configuration options. However, you might want to have additional configuration properties for a given instance without having to add a parameter. For example, to debug a server instance on a specific node, you might want to start the server with the `Ice.Trace.Network` property set; it would be inconvenient to have to add a parameter to the template just to set that property.

To cater for such scenarios, it is possible to specify additional properties for a server instance without modifying the template. You can define such properties in the `server-instance` element, for example:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application>
        ...
        <node name="Node2">
            <server-instance template="EncoderServerTemplate" index="2">
                <properties>
                    <property name="Ice.Trace.Network" value="2"/>
                </properties>
            </server-instance>
        </node>
    </application>
</icegrid>
```

This sets the `Ice.Trace.Network` property for a specific server.

Default Templates

The IceGrid registry can be configured to supply any number of default template descriptors for use in your applications. The configuration property `IceGrid.Registry.DefaultTemplates` specifies the path name of an XML file containing template definitions. One such template file is provided in the Ice distribution as `config/templates.xml`, which contains helpful templates for deploying Ice services such as IcePatch2 and Glacier2.

The template file must use the structure shown below:

XML

```
<icegrid>
    <application name="DefaultTemplates">
        <server-template id="EncoderServerTemplate">
            ...
        </server-template>
    </application>
</icegrid>
```

The name you give to the application is not important, and you may only define server and service templates within it. After configuring the registry to use this file, your default templates become available to every application that imports them.

The descriptor for each application indicates whether the default templates should be imported. (By default they are not imported.) If the templates are imported, they are essentially copied into the application descriptor and treated no differently than templates defined by the application itself. As a result, changes to the file containing default templates have no effect on existing application descriptors. In XML, the attribute `import-default-templates` determines whether the default templates are imported, as shown in the following example:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="Ripper" import-default-templates="true">
        ...
    </application>
</icegrid>
```

Using Templates with `icegridadmin`

The IceGrid administration tools allow you to inspect templates and instantiate new servers dynamically. First, let us ask `icegridadmin` to describe the server template we created earlier:

```
$ icegridadmin --Ice.Config=/opt/ripper/config
>>> server template describe Ripper EncoderServerTemplate
```

This command generates the following output:

```
server template `EncoderServerTemplate'
{
    parameters = `index exepath'
    server `EncoderServer${index}'
    {
        exe = `${exepath}'
        activation = `on-demand'
        properties
        {
            EncoderAdapter.Endpoints = `tcp'
        }
        adapter `EncoderAdapter'
        {
            id = `EncoderAdapter${index}'
            replica group id =
            endpoints = `tcp'
            register process = `false'
            server lifetime = `true'
        }
    }
}
```

Notice that the server ID is a parameterized value; it cannot be evaluated until the template is instantiated with values for its parameters.

Next, we can use `icegridadmin` to create an instance of the encoder server template on a new node:

```
>>> server template instantiate Ripper Node3 EncoderServerTemplate index=3
```


Copyright 2016, ZeroC, Inc.

The command requires that we identify the application, node and template, as well as supply any parameters needed by the template. The new server instance is permanently added to the registry's database, but if we intend to keep this configuration it is a good idea to update the XML description of our application to reflect these changes and avoid potential synchronization issues.

See Also

- Server Descriptor Element
- Server-Template Descriptor Element
- Server-Instance Descriptor Element
- icegridadmin Command Line Tool

Copyright 2016, ZeroC, Inc.

IceBox Integration with IceGrid

IceGrid makes it easy to configure an IceBox server with one or more services.

On this page:

- Deploying an IceBox Server
- Service Templates
- Advanced Service Templates

Deploying an IceBox Server

An IceBox server shares many of the same characteristics as other servers, but its special requirements necessitate a new descriptor. Unlike other servers, an IceBox server generally hosts multiple independent services, each requiring its own communicator instance and configuration file.

As an example, the following application deploys an IceBox server containing one service:

XML

```
<icegrid>
    <application name="IceBoxDemo">
        <node name="Node">
            <icebox id="IceBoxServer" exe="/opt/Ice/bin/icebox"
activation="on-demand">
                <service name="ServiceA" entry="servicea:create">
                    <adapter name="${service}" endpoints="tcp"/>
                </service>
            </icebox>
        </node>
    </application>
</icegrid>
```

It looks very similar to a server descriptor. The most significant difference is the service descriptor, which is constructed much like a server in that you can declare its attributes such as object adapters and configuration properties. The order in which services are defined determines the order in which they are loaded by the IceBox server.

The value of the adapter's `name` attribute needs additional explanation. The symbol `service` is one of the names reserved by IceGrid. In the context of a service descriptor, `${service}` is replaced with the service's name, and so the object adapter is also named `ServiceA`.

Service Templates

If you are familiar with templates in general, an IceBox service template is readily understandable:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="IceBoxApp">
        <service-template id="ServiceTemplate">
            <parameter name="name"/>
            <service name="${name}" entry="DemoService:create">
                <adapter name="${service}" endpoints="default"/>
                <property name="${service}.Identity" value="${server}-${service}"/>
            </service>
        </service-template>
        <node name="Node1">
            <icebox id="IceBoxServer" endpoints="default"
                exe="/opt/Ice/bin/icebox" activation="on-demand">
                <service-instance template="ServiceTemplate" name="Service1"/>
            </icebox>
        </node>
    </application>
</icegrid>
```

In this application, an IceBox server is deployed on a node and has one service instantiated from the service template. Of particular interest is the `property` descriptor, which uses another reserved name `server` to form the property value. When the template is instantiated by the service instance descriptor, the symbol `${server}` is replaced with the name of the enclosing server, so the property definition expands as follows:

```
Service1.Identity=IceBoxServer-Service1
```

As with server instances, you can specify additional properties for the service instance without modifying the template. These properties can be defined in the `service-instance` element, as shown below:

XML

```
<icegrid>
    <application name="IceBoxApp">
        ...
        <node name="Node1">
            <icebox id="IceBoxServer"endpoints="default"
                exe="/opt/Ice/bin/icebox" activation="on-demand">
                <service-instance template="ServiceTemplate" name="Service1">
                    <properties>
                        <property name="Ice.Trace.Network" value="1"/>
                    </properties>
                </service-instance>
            </icebox>
        </node>
    </application>
</icegrid>
```


Copyright 2016, ZeroC, Inc.

Advanced Service Templates

A more sophisticated use of templates involves instantiating a service template in a server template:

XML

```
<icegrid>
    <application name="IceBoxApp">
        <service-template id="ServiceTemplate">
            <parameter name="name"/>
            <service name="${name}" entry="DemoService:create">
                <adapter name="${service}" endpoints="default"/>
                <property name="${name}.Identity" value="${server}-${name}"/>
            </service>
        </service-template>
        <server-template id="ServerTemplate">
            <parameter name="id"/>
            <icebox id="${id}" endpoints="default"
                exe="/opt/Ice/bin/icebox" activation="on-demand">
                <service-instance template="ServiceTemplate" name="Service1"/>
            </icebox>
        </server-template>
        <node name="Node1">
            <server-instance template="ServerTemplate" id="IceBoxServer"/>
        </node>
    </application>
</icegrid>
```

This application is equivalent to our first example of service templates. Now, however, the process of deploying an identical server on several nodes has become much simpler.

If you need the ability to customize the configuration of a particular service instance, your server instance can define a property set that applies only to the desired service:

XML

```
<icegrid>
    <application name="IceBoxApp">
        <node name="Node1">
            <server-instance template="ServerTemplate" id="IceBoxServer">
                <properties service="Service1">
                    <property name="Ice.Trace.Network" value="1"/>
                </properties>
            </server-instance>
        </node>
    </application>
</icegrid>
```

As this example demonstrates, the `service` attribute of the property set denotes the name of the target service.

See Also

Copyright 2016, ZeroC, Inc.

- IceBox
- IceBox Descriptor Element
- Service Descriptor Element
- Service-Template Descriptor Element
- Server-Template Descriptor Element
- Properties Descriptor Element
- Using Descriptor Variables and Parameters

Copyright 2016, ZeroC, Inc.

Object Adapter Replication

As an implementation of an Ice location service, IceGrid supports object adapter replication. An application defines its replica groups and their participating object adapters using descriptors, and IceGrid generates the server configurations automatically.

On this page:

- Deploying a Replica Group
- Replica Group Membership
- Using Replica Groups in the Ripper Application
  - Adding a Replica Group to the Ripper Deployment
  - Using a Replica Group in the Ripper Client

Deploying a Replica Group

The descriptor that defines a replica group can optionally declare well-known objects as well as configure the group to determine its behavior during locate requests. Consider this example:

XML

```
<icegrid>
    <application name="ReplicaApp">
        <replica-group id="ReplicatedAdapter">
            <object identity="TheObject" type="::Demo::ObjectType"/>
        </replica-group>
        <node name="Node">
            <server id="ReplicaServer" activation="on-demand"
exe="/opt/replica/bin/server">
                <adapter name="TheAdapter" endpoints="default"
replica-group="ReplicatedAdapter"/>
            </server>
        </node>
    </application>
</icegrid>
```

The adapter's descriptor declares itself to be a member of the replica group `ReplicatedAdapter`, which must have been previously created by a replica group descriptor.

The replica group `ReplicatedAdapter` declares a well-known object so that an indirect proxy of the form `TheObject` is equivalent to the indirect proxy `TheObject@ReplicatedAdapter`. Since this trivial example defines only one adapter in the replica group, the proxy `TheOb ject` is also equivalent to `TheObject@TheAdapter`.

Replica Group Membership

An object adapter participates in a replica group by specifying the group's ID in the adapter's `ReplicaGroupId` configuration property. Identifying the replica group in the IceGrid descriptor for an object adapter causes the node to include the equivalent `ReplicaGroupId` prop erty in the configuration file it generates for the server.

By default, the IceGrid registry requires the membership of a replica group to be statically defined. When you create a descriptor for an object adapter that identifies a replica group, the registry adds that adapter to the group's list of valid members. During an adapter's activation, when it describes its endpoints to the registry, an adapter that also claims membership in a replica group is validated against the registry's internal list.

In a properly configured IceGrid application, this activity occurs without incident, but there are situations in which validation can fail. For example, adapter activation fails if an adapter's ID is changed without notifying the registry, such as by manually modifying the server configuration file that was generated by a node.

It is also possible for activation to fail when the IceGrid registry is being used solely as a location service, in which case descriptors have not been created and therefore the registry has no advance knowledge of the replica groups or their members. In this situation, adapter activation causes the server to receive `NotRegisteredException` unless the registry is configured to allow dynamic registration, which

Copyright 2016, ZeroC, Inc.

you can do by defining the following property:

```
IceGrid.Registry.DynamicRegistration=1
```

With this configuration, a replica group is created implicitly as soon as an adapter declares membership in it, and any adapter is allowed to participate.

The use of dynamic registration often leads to the accumulation of obsolete replica groups and adapters in the registry. The IceGrid administration tools allow you to inspect and clean up the registry's state.

Using Replica Groups in the Ripper Application

Replication is a perfect fit for the ripper application. The collection of encoder factory objects should be treated as a single logical object, and replication makes that possible.

Adding a Replica Group to the Ripper Deployment

Adding a replica group descriptor to our application is very straightforward:

XML

```
<icegrid>
    <application name="Ripper">
        <replica-group id="EncoderAdapters">
            <object identity="EncoderFactory" type="::Ripper::MP3EncoderFactory"/>
        </replica-group>
        <server-template id="EncoderServerTemplate">
            <parameter name="index"/>
            <parameter name="exepath" default="/opt/ripper/bin/server"/>
            <server id="EncoderServer${index}" exe="${exepath}"
activation="on-demand">
                <adapter name="EncoderAdapter" replica-group="EncoderAdapters"
                    endpoints="tcp"/>
            </server>
        </server-template>
        <node name="Node1">
            <server-instance template="EncoderServerTemplate" index="1"/>
        </node>
        <node name="Node2">
            <server-instance template="EncoderServerTemplate" index="2"/>
        </node>
    </application>
</icegrid>
```

The new descriptor adds the replica group called `EncoderAdapters` and registers a well-known object with the identity `EncoderFactory`. The adapter descriptor in the server template has been changed to declare its membership in the replica group.

Using a Replica Group in the Ripper Client

In comparison to the examples that demonstrated querying for well-known objects, the new version of our client has become much simpler:

Copyright 2016, ZeroC, Inc.

```
Ice::ObjectPrx obj = communicator->stringToProxy("EncoderFactory");
Ripper::MP3EncoderFactoryPrx factory =
Ripper::MP3EncoderFactoryPrx::checkedCast(obj);
Ripper::MP3EncoderPrx encoder = factory->createEncoder();
```

The client no longer needs to use the `IceGrid::Query` interface, but simply creates a proxy for a well-known object and lets the Ice run time transparently interact with the location service. In response to a locate request for `EncoderFactory`, the registry returns a proxy containing the endpoints of both object adapters. The Ice run time in the client selects one of the endpoints at random, meaning we have now lost some functionality compared to the prior example in which system load was considered when selecting an endpoint. We will learn how to rectify this situation in our discussion of load balancing.

See Also

- Terminology
- Replica-Group Descriptor Element
- Object Descriptor Element
- Well-Known Objects
- Load Balancing

Copyright 2016, ZeroC, Inc.

Load Balancing

Replication is an important IceGrid feature but, when combined with load balancing, replication becomes even more useful.

IceGrid nodes regularly report the system load of their hosts to the registry. The replica group's configuration determines whether the registry actually considers system load information while processing a locate request. Its configuration also specifies how many replicas to include in the registry's response.

IceGrid's load balancing capability assists the client in obtaining an initial set of endpoints for the purpose of establishing a connection. Once a client has established a connection, all subsequent requests on the proxy that initiated the connection are normally sent to the same server without further consultation with the registry. As a result, the registry's response to a locate request can only be viewed as a snapshot of the replicas at a particular moment. If system loads are important to the client, it must take steps to periodically contact the registry and update its endpoints.

On this page:

- Replica Group Load Balancing
- Load Balancing Types
- Using Load Balancing in the Ripper Application
- Interacting with Object Replicas
- Custom Load Balancing Strategies
  - Overview of Custom Load Balancing
    - The Registry Plug-in Facade Object
    - Implementing a Registry Plug-in
    - Installing a Registry Plug-in
    - Filter Implementation Techniques
  - Implementing a Custom Replica Group Filter
  - Implementing a Custom Type Filter

Replica Group Load Balancing

A replica group descriptor optionally contains a load balancing descriptor that determines how system loads are used in locate requests. The load balancing descriptor specifies the following information:

- Type
  Several load balancing types are supported.

- Sampling interval
  One of the load balancing types considers system load statistics, which are reported by each node at regular intervals. The replica group can specify a sampling interval of one, five, or fifteen minutes. Choosing a sampling interval requires balancing the need for up-to-date load information against the desire to minimize transient spikes. On Unix platforms, the node reports the system's load average for the selected interval, while on Windows the node reports the CPU utilization averaged over the interval.

- Number of replicas
  The replica group can instruct the registry to return the endpoints of one (the default) or more object adapters. If the specified number *N* is larger than one, the proxy returned in response to a locate request contains the endpoints of at most *N* object adapters. If *N* is 0, the proxy contains the endpoints of all the object adapters. The Ice run time in the client selects one of these endpoints at random when establishing a connection.

For example, the descriptor shown below uses adaptive load balancing to return the endpoints of the two least-loaded object adapters sampled with five-minute intervals:

XML

```
<replica-group id="ReplicatedAdapter">
    <load-balancing type="adaptive" load-sample="5" n-replicas="2"/>
</replica-group>
```

The type must be specified, but the remaining attributes are optional.

As of Ice 3.5, IceGrid ignores the object adapters of a disabled server when executing a locate request, meaning the client that initiated the locate request will not receive the endpoints for any of these object adapters.

Copyright 2016, ZeroC, Inc.

As of Ice 3.6, you can optionally use custom load balancing strategies by installing replica group filters.

Load Balancing Types

A replica group can select one of the following load balancing types:

- Random
  Random load balancing selects the requested number of object adapters at random. The registry does not consider system load for a replica group with this type.
- Adaptive
  Adaptive load balancing uses system load information to choose the least-loaded object adapters over the requested sampling interval. This is the only load balancing type that uses sampling intervals.
- Round Robin
  Round robin load balancing returns the least recently used object adapters. The registry does not consider system load for a replica group with this type. Note that the round-robin information is not shared between registry replicas; each replica maintains its own notion of the "least recently used" object adapters.
- Ordered
  Ordered load balancing selects the requested number of object adapters by priority. A priority can be set for each object adapter member of the replica group. If you define several object adapters with the same priority, IceGrid will order these object adapters according to their order of appearance in the descriptor.

Choosing the proper type of load balancing is highly dependent on the needs of client applications. Achieving the desired load balancing and fail-over behavior may also require the cooperation of your clients. To that end, it is very important that you understand how and when the Ice run time uses a locator to resolve indirect proxies.

Using Load Balancing in the Ripper Application

The only change we need to make to the ripper application is the addition of a load balancing descriptor:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="Ripper">
        <replica-group id="EncoderAdapters">
            <load-balancing type="adaptive"/>
            <object identity="EncoderFactory" type="::Ripper::MP3EncoderFactory"/>
        </replica-group>
        <server-template id="EncoderServerTemplate">
            <parameter name="index"/>
            <parameter name="exepath" default="/opt/ripper/bin/server"/>
            <server id="EncoderServer${index}" exe="${exepath}"
activation="on-demand">
                <adapter name="EncoderAdapter" replica-group="EncoderAdapters"
                    endpoints="tcp"/>
            </server>
        </server-template>
        <node name="Node1">
            <server-instance template="EncoderServerTemplate" index="1"/>
        </node>
        <node name="Node2">
            <server-instance template="EncoderServerTemplate" index="2"/>
        </node>
    </application>
</icegrid>
```

Using adaptive load balancing, we have regained the functionality we forfeited when we introduced replica groups. Namely, we now select the object adapter on the least-loaded node, and no changes are necessary in the client.

Interacting with Object Replicas

In some applications you may have a need for interacting directly with the replicas of an object. You might be tempted to call `ice_getEndpoints` on the proxy of a replicated object in an effort to obtain the endpoints of all replicas, but that is not the correct solution because the proxy is indirect and therefore contains no endpoints. The proper approach is to query well-known objects using the `findAllReplicas` operation.

Custom Load Balancing Strategies

As of Ice 3.6, the IceGrid registry allows you to plug in custom load balancing implementations that the registry invokes to filter its query results. Two kinds of filters are supported:

- Replica group filter
  The registry invokes a replica group filter each time a client requests the endpoints of a replica group or object adapter, as well as for calls to `findAllReplicas`. The registry passes information about the query that the filter can use in its implementation, including the list of object adapters participating in the replica group whose nodes are active at the time of the request. The object adapter list is initially ordered using the load balancing type configured for the replica group; the filter can modify this list however it chooses.

- Type filter
  The registry invokes a type filter for each query that a client issues to find a well-known object by type using the operations `findObjectByType`, `findAllObjectsByType`, and `findObjectByTypeOnLeastLoadedNode`. Included in the information passed to the filter is a list of proxies for the matching well-known objects; the filter implementation decides which of these proxies are returned to the client.

In the sections below we describe how to implement these filters.

Copyright 2016, ZeroC, Inc.

Overview of Custom Load Balancing

Filters are installed into the IceGrid registry using the standard Ice plug-in facility. The registry is implemented in C++, therefore the filters must be implemented in C++.

The Registry Plug-in Facade Object

During initialization, your plug-in will obtain a reference to a facade object with which it can register one or more filters. A filter typically retains a reference to this facade object because it offers a number of useful methods that the filter might need during its implementation. The `RegistryPluginFacade` class provides the following methods:

C++

```
namespace IceGrid {
class RegistryPluginFacade : virtual public Ice::LocalObject {
public:
    void addReplicaGroupFilter(const std::string& id, const
IceGrid::ReplicaGroupFilterPtr& filter);
    bool removeReplicaGroupFilter(const std::string& id, const
IceGrid::ReplicaGroupFilterPtr& filter);

    void addTypeFilter(const std::string& id, const IceGrid::TypeFilterPtr&
filter);
    bool removeTypeFilter(const std::string& id, const IceGrid::TypeFilterPtr&
filter);

    IceGrid::ApplicationInfo getApplicationInfo(const std::string& name) const;

    IceGrid::ServerInfo getServerInfo(const std::string& serverId) const;

    std::string getAdapterServer(const std::string& adapterId) const;
    std::string getAdapterApplication(const std::string& adapterId) const;
    std::string getAdapterNode(const std::string& adapterId) const;
    IceGrid::AdapterInfoSeq getAdapterInfo(const std::string& adapterId) const;
    std::string getPropertyForAdapter(const std::string& adapterId, const
std::string& name) const;

    IceGrid::ObjectInfo getObjectInfo(const Ice::Identity& id) const;

    IceGrid::NodeInfo getNodeInfo(const std::string& name) const;
    IceGrid::LoadInfo getNodeLoad(const std::string& name) const;
};
typedef ... RegistryPluginFacadePtr;

RegistryPluginFacadePtr getRegistryPluginFacade();
}
```

There are methods for adding and removing replica group and type filters, along with a number of methods for obtaining information about the deployment. As you can see, a great deal of information is available to a filter implementation for use in making its decisions. The data structures returned by these methods correspond directly to their XML descriptors; you can also review the Slice definitions of the IceGrid data types for more information.

The methods are described below:

- `void addReplicaGroupFilter(const std::string& id, const IceGrid::ReplicaGroupFilterPtr& filter)`
  Adds a replica group filter with the given identifier. The identifier must match the `filter` attribute of a `replica-group` descriptor. The registry maintains a list of filters for each identifier. If you register more than one filter with the same identifier, the registry

Copyright 2016, ZeroC, Inc.

invokes each one in turn, in the order of registration. To add a replica group filter for dynamically registered replica groups, you should use the empty string for the identifier.

- `bool removeReplicaGroupFilter(const std::string& id, const IceGrid::ReplicaGroupFilterPtr& filter)`
  Removes an existing replica group filter that matches the given identifier and smart pointer. Returns true if a match was found, false otherwise.
- `void addTypeFilter(const std::string& id, const IceGrid::TypeFilterPtr& filter)`
  Adds a type filter with the given identifier. The registry maintains a list of filters for each identifier. If you register more than one filter with the same identifier, the registry invokes each one in turn, in the order of registration.
- `bool removeTypeFilter(const std::string& id, const IceGrid::TypeFilterPtr& filter)`
  Removes an existing type filter that matches the given identifier and smart pointer. Returns true if a match was found, false otherwise.
- `IceGrid::ApplicationInfo getApplicationInfo(const std::string& name) const`
  Returns the descriptor for the application with the given name. Raises `IceGrid::ApplicationNotExistException` if no match is found.
- `IceGrid::ServerInfo getServerInfo(const std::string& serverId) const`
  Returns the descriptor for the server with the given identifier. Raises `IceGrid::ServerNotExistException` if no match is found.
- `std::string getAdapterServer(const std::string& adapterId) const`
  Returns the identifier of the server hosting the object adapter with the given adapter identifier. Returns an empty string if no match is found.
- `std::string getAdapterApplication(const std::string& adapterId) const`
  Returns the name of the application containing the given object adapter identifier. Returns an empty string if no match is found.
- `std::string getAdapterNode(const std::string& adapterId) const`
  Returns the name of the node containing the given object adapter identifier. Returns an empty string if no match is found.
- `IceGrid::AdapterInfoSeq getAdapterInfo(const std::string& adapterId) const`
  If `adapterId` is a replica group identifier, this method returns a sequence of adapter descriptors for all of the replicas. Otherwise, this method returns a sequence containing one descriptor for the given object adapter. Raises `IceGrid::AdapterNotExistException` if no match is found.
- `std::string getPropertyForAdapter(const std::string& adapterId, const std::string& name) const`
  Obtains the value of a server configuration property with the key `name` for the server hosting the object adapter identified by `adapterId`. Returns an empty string if no match is found.
- `IceGrid::ObjectInfo getObjectInfo(const Ice::Identity& id) const`
  Returns information about an object with the given identity. Raises `IceGrid::ObjectNotRegisteredException` if no match is found.
- `IceGrid::NodeInfo getNodeInfo(const std::string& name) const`
  Returns the descriptor for the node with the given name. Raises `IceGrid::NodeNotExistException` if no match is found.
- `IceGrid::LoadInfo getNodeLoad(const std::string& name) const`
  Returns load information for the node with the given name. Raises `IceGrid::NodeNotExistException` if no match is found.

Methods that require access to deployment information raise `IceGrid::RegistryUnreachableException` if called before the registry has fully initialized itself.

Implementing a Registry Plug-in

The API for creating an Ice plug-in using C++ requires a factory function with external linkage, along with a class that implements the `Ice::Plugin` local Slice interface. We can use the code from the example in `demo/IceGrid/customLoadBalancing` to illustrate these points. First, the factory function instantiates and returns the plug-in:

Copyright 2016, ZeroC, Inc.

C++

```
extern "C" {
ICE_DECLSPEC_EXPORT Ice::Plugin*
createRegistryPlugin(const Ice::CommunicatorPtr& communicator, const string&, const
Ice::StringSeq&)
{
    return new RegistryPluginI(communicator);
}
}
```

The plug-in class is straightforward:

C++

```
class RegistryPluginI : public Ice::Plugin {
public:
    RegistryPluginI(const Ice::CommunicatorPtr& communicator) :
_communicator(communicator) {}

    virtual void initialize()
    {
        IceGrid::RegistryPluginFacadePtr facade =
IceGrid::getRegistryPluginFacade();
        if(facade)
        {
            facade->addReplicaGroupFilter("filterByCurrency", new
ReplicaGroupFilterI(facade));
        }
    }

    virtual void destroy() {}

private:
    const Ice::CommunicatorPtr _communicator;
};
```

The `initialize` method calls `getRegistryPluginFacade` to obtain a smart pointer for the registry's facade object. The plug-in uses this object to install a replica group filter.

We describe filter implementations in more detail below.

Installing a Registry Plug-in

Continuing with our example from `demo/IceGrid/customLoadBalancing`, we use the following property to install our plug-in in the IceGrid registry:

```
Ice.Plugin.RegistryPlugin=RegistryPlugin:createRegistryPlugin
```

The `Ice.Plugin` property must be defined in the registry's configuration file.

Copyright 2016, ZeroC, Inc.

Make sure to configure all of the replicas to load the same registry plug-ins, otherwise a client could get different behavior depending on which replica it's currently using.

Filter Implementation Techniques

A filter may require client-specific information in order to assemble its list of results. We recommend using request contexts for this purpose. Briefly, a request context is a dictionary of key/value string pairs that a client can configure and send along as "out of band" metadata accompanying a request. Ice provides several ways for a client to establish a request context:

- Implicit - provides a default request context for every request on all proxies
- Per-proxy - configures a default request context for every request on a particular proxy
- Explicit - specifies a request context at the time of invocation, overriding any default context

Our goal here is to transfer information from a client to a filter. Consequently, we don't recommend using implicit contexts because the context will add overhead to *every* invocation the client makes, not just the ones that involve the filter. To decide whether to use per-proxy or explicit contexts, we first must understand the circumstances in which each kind of filter receives a request context:

- Replica group filter

  Most invocations involving a replica group filter occur when the Ice run time in a client issues requests on a registry. This internal activity is triggered by the client's invocation on a proxy, but Ice doesn't use the client's proxy or the request context that the client may have configured for its proxy or passed explicitly to the invocation. Rather, the Ice run time uses the locator that the client configured for its proxy or communicator. As a result, the request contexts passed to a replica group filter are those configured for the *locator* proxy. (The only exception is when a client invokes `findAllReplicas` directly on the registry via its `Query` interface; refer to the discussion of type filters below for more details.)

  There are several ways you can configure a request context for the locator proxy. If a client configures its locator proxy statically using properties, the simplest solution is to add `Context` properties to the client's configuration, such as:

  ```
  Ice.Default.Locator=...
  Ice.Default.Locator.Context.someKey=someValue
  ```

  If you need to define the context at run time, you can obtain the locator proxy by calling `getDefaultLocator` on the communicator, create a new locator proxy with the desired context by calling `ice_context`, and then replace the locator proxy by calling `setDefaultLocator`.

  Both of these approaches are examples of per-proxy request contexts.

- Type filter

  Invocations involving type filters occur when a client invokes directly on the registry using its `Query` interface. To use per-proxy request contexts, configure the `Query` proxy as necessary. Explicit request contexts can also be used when querying the registry.

So far we've discussed how client-specific information can be passed to a filter, but what if the filter needs to obtain more information about the object adapters (in the case of a replica group filter) or objects (in the case of a type filter) in order to perform its duties? This is where the registry's facade object comes in handy, as with it the filter can retrieve information about the deployment. For example, a replica group filter can use server-specific properties as a form of metadata. The registry supplies the filter with a list of object adapter identifiers; each object adapter is hosted by a server, and the filter can look up property values for that server using the facade. The sample filter in `demo/IceGrid/customLoadBalancing` uses this technique, and we describe it in more detail below.

Implementing a Custom Replica Group Filter

A replica group filter must define a subclass of `ReplicaGroupFilter`:

Copyright 2016, ZeroC, Inc.

C++

```
namespace IceGrid {
class ReplicaGroupFilter : virtual public Ice::LocalObject {
public:
    virtual Ice::StringSeq filter(const std::string& id,
                                  const Ice::StringSeq& adapters,
                                  const Ice::ConnectionPtr& connection,
                                  const Ice::Context& context);
};
}
```

The registry passes the following arguments to the `filter` method:

- `id`
  The replica group identifier involved in this request.
- `adapters`
  A sequence of object adapter identifiers denoting the object adapters participating in the replica group whose nodes are active at the time of the request. The object adapter list is initially ordered using the load balancing type configured for the replica group.
- `connection`
  The incoming connection from the client to the registry.
- `context`
  The request context that accompanied the request.

The implementation returns a sequence containing zero or more object adapter identifiers. The registry may truncate this list if the filter supplies more object adapters than the `n-replicas` value configured for the replica group's load balancing policy, but the registry will not change the ordering.

The `filter` implementation must not block.

The C++ example in `demo/IceGrid/customLoadBalancing` uses a replica group filter to select only those object adapters that support the currency requested by the client. The replica group's descriptor specifies the filter:

```
<replica-group id="ReplicatedPricingAdapter" filter="filterByCurrency">
  <load-balancing type="random"/>
  <object identity="pricing" type="::Demo::PricingEngine"/>
</replica-group>
```

Notice that the filter identifier `filterByCurrency` matches that used when the plug-in registered the filter.

In this example, the client uses a request context to indicate the desired currency. The context is configured on the locator proxy in the client's configuration file:

```
Ice.Default.Locator=...
Ice.Default.Locator.Context.currency=USD
```

Here we use the `Context` proxy property to statically assign a request context to the locator proxy, which means every invocation on the locator proxy includes the key/value pair `currency/USD`.

In addition to configuring the replica group filter, the deployment descriptor plays another important role here by defining server-specific properties that the filter uses in its implementation:

Copyright 2016, ZeroC, Inc.

```
<server-template id="PricingServer">
  <parameter name="index"/>
  <parameter name="currencies"/>
  <server id="PricingServer-${index}" exe="./server" activation="on-demand">
    <adapter name="Pricing" endpoints="tcp -h localhost"
replica-group="ReplicatedPricingAdapter"/>
    <property name="Identity" value="pricing"/>
    <property name="Currencies" value="${currencies}"/>
    ...
  </server>
</server-template>


<node name="node1">
  <server-instance template="PricingServer" index="1" currencies="EUR, GBP, JPY"/>
  <server-instance template="PricingServer" index="2" currencies="USD, GBP, AUD"/>
  <server-instance template="PricingServer" index="3" currencies="EUR, USD, INR"/>
  <server-instance template="PricingServer" index="4" currencies="JPY, GBP, AUD"/>
  <server-instance template="PricingServer" index="5" currencies="GBP, AUD"/>
  <server-instance template="PricingServer" index="6" currencies="EUR, USD"/>
</node>
```

As a convenience, the descriptor file uses a server template to define several servers, with each server supporting a specific set of currencies. The template adds to the property set of each server a property named `Currencies` representing the currencies that the server supports.

Finally, you can see how all of this ties together in the filter implementation:

Copyright 2016, ZeroC, Inc.

```
Ice::StringSeq
ReplicaGroupFilterI::filter(const string& replicaGroupId,
                            const Ice::StringSeq& adapters,
                            const Ice::ConnectionPtr& connection,
                            const Ice::Context& ctx)
{
    Ice::Context::const_iterator p = ctx.find("currency");
    if(p == ctx.end())
    {
        return adapters;
    }

    string currency = p->second;

    Ice::StringSeq filteredAdapters;
    for(Ice::StringSeq::const_iterator p = adapters.begin(); p != adapters.end();
++p)
    {
        if(_facade->getPropertyForAdapter(*p, "Currencies").find(currency) !=
string::npos)
        {
            filteredAdapters.push_back(*p);
        }
    }
    return filteredAdapters;
```

The code first checks the request context for a value associated with the key `currency` and returns the adapter list unmodified if no entry is found. Next, the method builds a new list of adapters by iterating over the adapter list in its existing order and calling `getPropertyForAdapter` on the facade for each adapter. This method uses the registry's deployment information to find the server hosting the given adapter and then searches the server's configuration properties for one matching the given property name. If the `Currencies` property contains the client's specified currency, the adapter is added to the list that is eventually returned by the filter.

As this example demonstrates, request contexts are a convenient way to supply a filter with client-specific information, and server properties can serve as a simple database when a filter needs to tailor its results based on server attributes.

Implementing a Custom Type Filter

A replica group filter must define a subclass of `TypeFilter`:

C++

```
namespace IceGrid {
class TypeFilter : virtual public Ice::LocalObject {
public:
    virtual Ice::ObjectProxySeq filter(const std::string& type,
                                       const Ice::ObjectProxySeq& proxies,
                                       const Ice::ConnectionPtr& connection,
                                       const Ice::Context& context);
};
}
```


Copyright 2016, ZeroC, Inc.

The registry passes the following arguments to the `filter` method:

- `type`
  The type identifier involved in this request.
- `proxies`
  A sequence of proxies denoting the objects that matched the type.
- `connection`
  The incoming connection from the client to the registry.
- `context`
  The request context that accompanied the request.

The implementation returns a sequence containing zero or more proxies.

The `filter` implementation must not block.

Refer to the previous section for more information on implementing a filter.

See Also

- Object Adapter Replication
- Connection Establishment
- Replica-Group Descriptor Element
- Load-Balancing Descriptor Element
- Well-Known Objects
- IceGrid Troubleshooting

Copyright 2016, ZeroC, Inc.

Resource Allocation using IceGrid Sessions

IceGrid provides a resource allocation facility that coordinates access to the objects and servers of an IceGrid application. To allocate a resource for exclusive use, a client must first establish a session by authenticating itself with the IceGrid registry or a Glacier2 router, after which the client may reserve objects and servers that the application indicates are allocatable. The client should release the resource when it is no longer needed, otherwise IceGrid reclaims it when the client's session terminates or expires due to inactivity.

An allocatable server offers at least one allocatable object. The server is considered to be allocated when its first allocatable object is claimed, and is not released until all of its allocated objects are released. While the server is allocated by a client, no other clients can allocate its objects.

On this page:

- Creating an IceGrid Session
- Controlling Access to IceGrid Sessions
- Allocating Objects with an IceGrid Session
- Allocating Servers with an IceGrid Session
- Security Considerations for Allocated Resources
- Deploying Allocatable Resources
- Using Resource Allocation in the Ripper Application

Creating an IceGrid Session

A client must create an IceGrid session before it can allocate objects. If you have configured a Glacier2 router to use IceGrid's session managers, the client's router session satisfies this requirement.

In the absence of Glacier2, an IceGrid client invokes `createSession` or `createSessionFromSecureConnection` on IceGrid's `Registry` interface to create a session:

Slice

```
module IceGrid {
    exception PermissionDeniedException {
        string reason;
    };

    interface Session extends Glacier2::Session {
        idempotent void keepAlive();
        ...
    };

    interface Registry {
        Session* createSession(string userId, string password)
            throws PermissionDeniedException;

        Session* createSessionFromSecureConnection()
            throws PermissionDeniedException;

        idempotent int getSessionTimeout();

        idempotent int getACMTimeout();

        ...
    };
};
```

The `createSession` operation expects a username and password and returns a session proxy if the client is allowed to create a session.

Copyright 2016, ZeroC, Inc.

By default, IceGrid does not allow the creation of sessions. You must define the registry property `IceGrid.Registry.PermissionsVerifier` with the proxy of a permissions verifier object to enable session creation with `createSession`.

The `createSessionFromSecureConnection` operation does not require a username and password because it uses the credentials supplied by an SSL connection to authenticate the client. As with `createSession`, you must enable session creation by configuring the proxy of a permissions verifier object so that clients can use `createSessionFromSecureConnection` to create a session. In this case, the property is `IceGrid.Registry.SSLPermissionsVerifier`.

To create a session, the client obtains the registry proxy by converting the well-known proxy string `"IceGrid/Registry"` to a proxy object with the communicator, downcasts the proxy to the `IceGrid::Registry` interface, and invokes one of the operations. The sample code below demonstrates how to do it in C++; the code will look very similar in other language mappings.

C++

```
Ice::ObjectPrx base = communicator->stringToProxy("IceGrid/Registry");
IceGrid::RegistryPrx registry = IceGrid::RegistryPrx::checkedCast(base);
string username = ...;
string password = ...;
IceGrid::SessionPrx session;
try {
    session = registry->createSession(username, password);
} catch (const IceGrid::PermissionDeniedException & ex) {
    cout << "permission denied:\n" << ex.reason << endl;
}
```

The identity of the registry object may change based on its configuration settings.

After creating the session, the client must keep it alive to prevent it from expiring. You have two options for keeping a session alive:

1. Call `Registry::getSessionTimeout` and periodically invoke `Session::keepAlive` every *timeout* (or less) seconds
2. Call `Registry::getACMTimeout` and configure ACM settings on the connection

We recommend using the second approach, which you can implement as follows:

C++

```
int acmTimeout = registry->getACMTimeout();
if (acmTimeout > 0) {
    Ice::ConnectionPtr conn = session->ice_getCachedConnection();
    conn->setACM(acmTimeout, IceUtil::None, Ice::HeartbeatAlways);
}
```

Enabling heartbeats on the connection causes Ice to automatically send a heartbeat message at regular intervals determined by the given timeout value. The server ignores these messages, but they serve the purpose of keeping the session alive.

The ACM features shown here were introduced in Ice 3.6.

If a session times out, or if the client explicitly terminates the session by invoking its `destroy` operation, IceGrid automatically releases all objects allocated using that session.

Controlling Access to IceGrid Sessions

Copyright 2016, ZeroC, Inc.

As described above, you must configure the IceGrid registry with the proxy of at least one permissions verifier object to enable session creation:

- `IceGrid.Registry.PermissionsVerifier`
  This property supplies the proxy of an object that implements the interface `Glacier2::PermissionsVerifier`. Defining this property allows clients to create sessions using `createSession`.
- `IceGrid.Registry.SSLPermissionsVerifier`
  This property supplies the proxy of an object that implements the interface `Glacier2::SSLPermissionsVerifier`. Defining this property allows clients to create sessions using `createSessionFromSecureConnection`.

IceGrid supplies built-in permissions verifier objects:

- A null permissions verifier for TCP/IP. This object accepts any username and password and should only be used in a secure environment where no access control is necessary. You select this verifier object by defining the following configuration property:

  ```
  IceGrid.Registry.PermissionsVerifier=<instance-name>/NullPermissionsVerifier
  ```

  Note that you have to substitute the correct instance name for the object identity category.
- A null permissions verifier for SSL, analogous to the one for TCP/IP. You select this verifier object by defining the following configuration property:

  ```
  IceGrid.Registry.SSLPermissionsVerifier=<instance-name>/NullSSLPermissionsVerifier
  ```
- A file-based permissions verifier. This object uses an access control list in a file that contains username-password pairs. The format of the password file is the same as the format of Glacier2 password files. You enable this verifier implementation by defining the configuration property `IceGrid.Registry.CryptPasswords` with the pathname of the password file. Note that this property is ignored if you specify the proxy of a permissions verifier object using `IceGrid.Registry.PermissionsVerifier`.

You can also implement your own permissions verifier object.

Allocating Objects with an IceGrid Session

A client allocates objects using the session proxy returned from `createSession` or `createSessionFromSecureConnection`. The proxy supports the `Session` interface shown below:

Copyright 2016, ZeroC, Inc.

Slice

```
module IceGrid {
    exception ObjectNotRegisteredException {
        Ice::Identity id;
    };

    exception AllocationException {
        string reason;
    };

    exception AllocationTimeoutException
        extends AllocationException {
    };

    interface Session extends Glacier2::Session {

        idempotent void keepAlive();

        Object* allocateObjectById(Ice::Identity id)
            throws ObjectNotRegisteredException,
                   AllocationException;

        Object* allocateObjectByType(string type)
            throws AllocationException;

        void releaseObject(Ice::Identity id)
            throws ObjectNotRegisteredException,
                   AllocationException;

        idempotent void setAllocationTimeout(int timeout);
    };
};
```

The client is responsible for keeping the session alive by periodically invoking `keepAlive`, as discussed earlier.

The `allocateObjectById` operation allocates and returns the proxy for the allocatable object with the given identity. If no allocatable object with the given identity is registered, the client receives `ObjectNotRegisteredException`. If the object cannot be allocated, the client receives `AllocationException`. An allocation attempt can fail for the following reasons:

- the object is already allocated by the session
- the object is allocated by another session and did not become available during the configured allocation timeout period
- the session was destroyed.

The `allocateObjectByType` operation allocates and returns a proxy for an allocatable object registered with the given type. If more than one allocatable object is registered with the given type, the registry selects one at random. The client receives `AllocationException` if no objects with the given type could be allocated. An allocation attempt can fail for the following reasons:

- no objects are registered with the given type
- all objects with the given type are already allocated (either by this session or other sessions) and none became available during the configured allocation timeout period
- the session was destroyed.

The `releaseObject` operation releases an object allocated by the session. The client receives `ObjectNotRegisteredException` if no allocatable object is registered with the given identity and `AllocationException` if the object is not allocated by the session. Upon session destruction, IceGrid automatically releases all allocated objects.

Copyright 2016, ZeroC, Inc.

The `setAllocationTimeout` operation configures the timeout used by the allocation operations. If no allocatable objects are available when the client invokes `allocateObjectById` or `allocateObjectByType`, IceGrid waits for the specified timeout period for an allocatable object to become available. If the timeout expires, the client receives `AllocationTimeoutException`.

Allocating Servers with an IceGrid Session

A client does not need to explicitly allocate a server. If a server is allocatable, IceGrid implicitly allocates it to the first client that claims one of the server's allocatable objects. Likewise, IceGrid releases the server when all of its allocatable objects are released.

Server allocation is useful in two situations:

- Only allocatable servers can use the `session` activation mode, in which the server is activated on demand when allocated by a client and deactivated upon release.
- An allocatable server can be secured with IceSSL or Glacier2 so that its objects can only be invoked by the client that allocated it.

Security Considerations for Allocated Resources

IceGrid's resource allocation facility allows clients to coordinate access to objects and servers but does not place any restrictions on client invocations to allocated objects; any client that has a proxy for an allocated object could conceivably invoke an operation on it. IceGrid assumes that clients are cooperating with each other and respecting allocation semantics.

To prevent unauthorized clients from invoking operations on an allocated object or server, you can use IceSSL or Glacier2:

- Using IceSSL, you can secure access to a server or a particular object adapter with the properties `IceSSL.TrustOnly.Server` or `IceSSL.TrustOnly.Server.AdapterName`. For example, if you configure a server with the session activation mode, you can set one of the `IceSSL.TrustOnly` properties to the `${session.id}` variable, which is substituted with the session ID when the server is activated for the session. If the IceGrid session was created from a secure connection, the session ID will be the distinguished name associated with the secure connection, which effectively restricts access to the server or one of its adapters to the client that established the session with IceGrid.
- With Glacier2, you can secure access to an allocated object or the object adapters of an allocated server with the Glacier2 filtering mechanism. By default, IceGrid sessions created with a Glacier2 router are automatically given access to allocated objects, allocatable objects, certain well-known objects, and the object adapters of allocated servers.

Deploying Allocatable Resources

Allocatable objects are registered using a descriptor that is similar to well-known object descriptors. Allocatable objects cannot be replicated and therefore can only be specified within an object adapter descriptor.

Servers can be specified as allocatable by setting the server descriptor's `allocatable` attribute.

As an example, the following application defines an allocatable server and an allocatable object:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="Ripper">
        <node name="Node1">
            <server id="EncoderServer"
                exe="/opt/ripper/bin/server"
                activation="on-demand"
                allocatable="true">
                <adapter name="EncoderAdapter" id="EncoderAdapter" endpoints="tcp">

                    <allocatable identity="EncoderFactory" type="::Ripper::MP3Encod
erFactory"/>
                </adapter>
            </server>
        </node>
    </application>
</icegrid>
```

Using Resource Allocation in the Ripper Application

We can use the allocation facility in our MP3 encoder factory to coordinate access to the MP3 encoder factories. First we need to modify the descriptors to define an allocatable object:

XML

```
<icegrid>
    <application name="Ripper">
        <server-template id="EncoderServerTemplate">
            <parameter name="index"/>
            <server id="EncoderServer${index}"
                exe="/opt/ripper/bin/server"
                activation="on-demand">
                <adapter name="EncoderAdapter"  endpoints="tcp">
                    <allocatable identity="EncoderFactory${index}"
                        type="::Ripper::MP3EncoderFactory"/>
                </adapter>
            </server>
        </server-template>
        <node name="Node1">
            <server-instance template="EncoderServerTemplate" index="1"/>
        </node>
        <node name="Node2">
            <server-instance template="EncoderServerTemplate" index="2"/>
        </node>
    </application>
</icegrid>
```

Next, the client needs to create a session and allocate a factory:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx obj = session->allocateObjectByType(Ripper::MP3EncoderFactory::ice_s
taticId());
try {
    Ripper::MP3EncoderPrx encoder = factory->createEncoder();
    // Use the encoder to encode a file ...
}
catch (const Ice::LocalException & ex) {
    // There was a problem with the encoding, we catch the
    // exception to make sure we release the factory.
}
session->releaseObject(obj->ice_getIdentity());
```

It is important to release an allocated object when it is no longer needed so that other clients may use it. If you forget to release an object, it remains allocated until the session is destroyed.

See Also

- Getting Started with Glacier2
- IceSSL
- Well-Known Registry Objects
- Securing a Glacier2 Router
- Object Descriptor Element
- Allocatable Descriptor Element
- IceGrid.*
- IceSSL.*

Copyright 2016, ZeroC, Inc.

Registry Replication

The failure of an IceGrid registry or registry host can have serious consequences. A client can continue to use an existing connection to a server without interruption, but any activity that requires interaction with the registry is vulnerable to a single point of failure. As a result, the IceGrid registry supports replication using a master-slave configuration to provide high availability for applications that require it.

On this page:

- Registry Replication Architecture
- Capabilities of a Registry Replica
  - Locate Requests
  - Server Activation
  - Queries
  - Allocation
  - Administration
  - Glacier2 Support
- Configuring Registry Replication
  - Replicas
  - Clients
  - Nodes
  - Diagnostics
- Using Registry Replication with External Load Balancing

Registry Replication Architecture

In IceGrid's registry replication architecture, there is one master replica and any number of slave replicas. The master synchronizes its deployment information with the slaves so that any replica is capable of responding to locate requests, managing nodes, and starting servers on demand. Should the master registry or its host fail, properly configured clients transparently fail over to one of the slaves.

Each replica has a unique name. The name `Master` is reserved for the master replica, while replicas can use any name that can legally appear in an object identity.

The figure below illustrates the underlying concepts of registry replication:

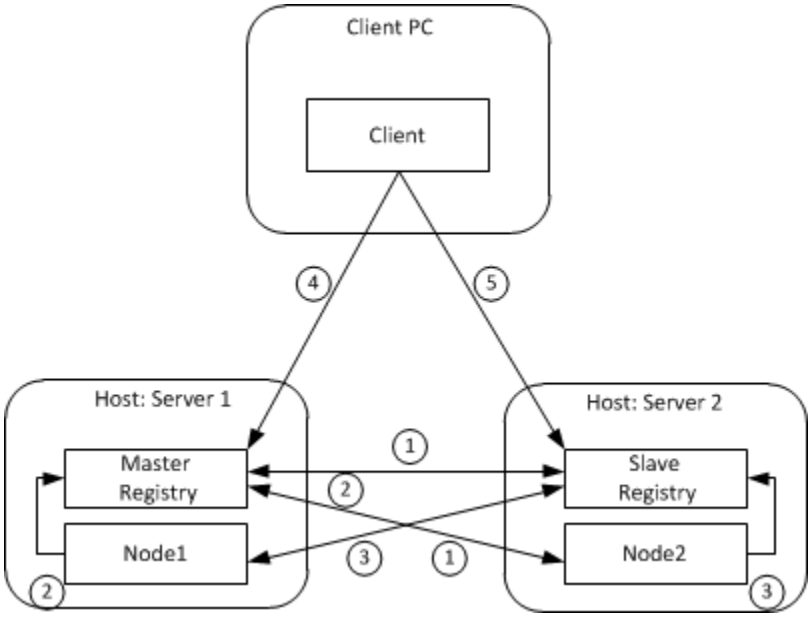

Overview of registry replication.

1. The slave replica contacts the master replica at startup and synchronizes its databases. Any subsequent modifications to the deployed applications are made via the master replica, which distributes them to all active slaves.
2. The nodes contact the master replica at startup to notify it about their availability.
3. The master replica provides a list of slave replicas to the nodes so that the nodes can also notify the slaves.
4. The client's configuration determines which replica it contacts initially. In this example, it contacts the master replica.
5. In the case of a failure, the client automatically fails over to the slave. If the master registry's host has failed, then `Node1` and any servers that were active on this host also become unavailable. The use of object adapter replication allows the client to transparently reestablish communication with a server on `Node2`.

Copyright 2016, ZeroC, Inc.

Capabilities of a Registry Replica

A master registry replica has a number of responsibilities, only some of which are supported by slaves. The master replica knows all of its slaves, but the slaves are not in contact with one another. If the master replica fails, the slaves can perform several vital functions that should keep most applications running without interruption. Eventually, however, a new master replica must be started to restore full registry functionality. For a slave replica to become the master, the slave must be restarted.

Locate Requests

One of the most important functions of a registry replica is responding to locate requests from clients, and every replica has the capability to service these requests. Slaves synchronize their databases with the master so that they have all of the information necessary to transform object identities, object adapter identifiers, and replica group identifiers into an appropriate set of endpoints.

Server Activation

Nodes establish sessions with each active registry replica so that any of the replicas are capable of activating a server on behalf of a client.

Queries

Replicating the registry also replicates the object that supports the `IceGrid::Query` interface used to query well-known objects. A client that resolves the `IceGrid/Query` object identity receives the endpoints of all active replicas, any of which can execute the client's requests.

Allocation

A client that needs to allocate a resource must establish a session with the master replica.

Administration

The state of an IceGrid registry is accessible via the `IceGrid::Admin` interface or (more commonly) using an administrative tool that encapsulates this interface. Modifications to the registry's state, such as deploying or updating an application, can only be done using the master replica. Administrative access to slave replicas is allowed but restricted to read-only operations. The administrative utilities provide mechanisms for you to select a particular replica to contact.

For programmatic access to a replica's administrative interface, the `IceGrid/Registry` identity corresponds to the master replica and the identity `IceGrid/Registry-`*`name`* corresponds to the slave with the given name.

Glacier2 Support

The registry implements the session manager interfaces required for integration with a Glacier2 router. The master replica supports the object identities `IceGrid/SessionManager` and `IceGrid/AdminSessionManager`. The slave replicas offer support for read-only administrative sessions using the object identity `IceGrid/AdminSessionManager-`*`name`*.

Configuring Registry Replication

Incorporating registry replication into an application is primarily accomplished by modifying your IceGrid configuration settings.

Replicas

Each replica must specify a unique name in its configuration property `IceGrid.Registry.ReplicaName`. The default value of this property is `Master`, therefore the master replica can omit this property if desired.

At startup, a slave replica attempts to register itself with its master in order to synchronize its databases and obtain the list of active nodes. The slave uses the proxy supplied by the `Ice.Default.Locator` property to find the master. At a minimum, this proxy should contain the endpoint of a replica that is connected to the master.

For better reliability if a failure occurs, we recommend that you also include the endpoints of all slave replicas in the `Ice.Default.Locator` property. There is no harm in adding the slave's own endpoints to the proxy in `Ice.Default.Locator`; in fact, it makes configuration simpler because all of the slaves can share the same property definition. Although slaves do not communicate with each other, it is possible for one of the slaves to be promoted to the master, therefore supplying the endpoints of all slaves minimizes the chance of a communication failure.

Shown below is an example of the configuration properties for a master replica:

Copyright 2016, ZeroC, Inc.

```
IceGrid.InstanceName=DemoIceGrid
IceGrid.Registry.Client.Endpoints=default -p 12000
IceGrid.Registry.Server.Endpoints=default
IceGrid.Registry.Internal.Endpoints=default
IceGrid.Registry.Data=db/master
...
```

You can configure a slave replica to use this master with the following settings:

```
Ice.Default.Locator=DemoIceGrid/Locator:default -p 12000
IceGrid.Registry.Client.Endpoints=default -p 12001
IceGrid.Registry.Server.Endpoints=default
IceGrid.Registry.Internal.Endpoints=default
IceGrid.Registry.Data=db/replica1
IceGrid.Registry.ReplicaName=Replica1
...
```


Configuring IceLocatorDiscovery in the replicas allows them to discover the master at run time without the need to define `Ice.Default.Locator`.

Clients

The endpoints contained in the `Ice.Default.Locator` property determine which registry replicas the client can use when issuing locate requests. If high availability is important, this property should include the endpoints of at least two (and preferably all) replicas. Not only does this increase the reliability of the client, it also distributes the work load of responding to locate requests among all of the replicas.

Continuing the example from the previous section, you can configure a client with the `Ice.Default.Locator` property as shown below:

```
Ice.Default.Locator=DemoIceGrid/Locator:default -p 12000:default -p 12001
```


Configuring IceLocatorDiscovery in a client allows it to discover the replicas at run time without the need to define `Ice.Default.Locator`.

Nodes

As with slave replicas and clients, an IceGrid node should be configured with an `Ice.Default.Locator` property that contains the endpoint of at least one replica and preferably all the replicas. A node needs to notify each of the registry replicas about its presence, thereby enabling the replicas to activate its servers and obtain the endpoints of its object adapters. Only the master replica knows the list of active replica slaves so it's important that the node connects to the master on startup to retrieve the list of all the active replicas. If the master is down when the node starts, the node will try to obtain the list of the registry replicas from the replicas specified in its `Ice.Default.Locator` proxy.

The following properties demonstrate how to configure a node with a replicated registry:

Copyright 2016, ZeroC, Inc.

```
Ice.Default.Locator=DemoIceGrid/Locator:default -p 12000:default -p 12001
IceGrid.Node.Name=node1
IceGrid.Node.Endpoints=default
IceGrid.Node.Data=db/node1
```

Configuring IceLocatorDiscovery in a node allows it to discover the replicas at run time without the need to define `Ice.Default.Locator`.

Diagnostics

You can use several configuration properties to enable trace messages that may help in diagnosing registry replication issues:

- `IceGrid.Registry.Trace.Replica`
  Displays information about the sessions established between master and slave replicas.
- `IceGrid.Registry.Trace.Node`
  `IceGrid.Node.Trace.Replica`
  Displays information about the sessions established between replicas and nodes.

Using Registry Replication with External Load Balancing

As explained earlier, we recommend including the endpoints of all replicas in the `Ice.Default.Locator` property. However, doing so might not always be convenient in large deployments, as it can require modifying many configuration files whenever you add or remove a registry replica. There are two ways to simplify your configuration:

- **Use a DNS name bound to multiple address (A) records**
  Configure the DNS name to point to the IP addresses of each of the registry replicas. You can then define the `Ice.Default.Locator` property with a single endpoint that embeds the DNS name. The Ice run time in the client randomly picks one of these IP addresses when it resolves the DNS name. All replicas must use the same port.
- **Use a TCP load balancer**
  Configure the load balancer to redirect traffic to the registry replicas and define the `Ice.Default.Locator` property with an endpoint that embeds the IP address and port of the TCP load balancer. The Ice run time in the client connects to the load balancer and the load balancer forwards the traffic to an active registry replica.

For maximum reliability, the DNS server or TCP load balancer should also be replicated.

When using such a configuration for the `Ice.Default.Locator` property of registry slaves, special care needs to be taken when initially starting the IceGrid registry replicas. Since the `Ice.Default.Locator` property no longer includes the endpoint of the master replica, a slave replica won't be able to contact the master if it's started when no master is running. You must first start the master and then the slaves to prevent this from happening. Once a replica slave connects successfully to the master, it saves the endpoint of the master and the other slaves in its database, so this is really only an issue when starting a slave with an empty database.

See Also

- Well-Known Objects
- icegridadmin Command Line Tool
- Glacier2 Integration with IceGrid
- IceGrid.*

Copyright 2016, ZeroC, Inc.

Application Distribution

On this page:

- Using IcePatch2 to Distribute Applications
- Deploying an IcePatch2 Server
  - Patching Considerations
  - IcePatch2 Server Template
- Adding Distribution to a Deployment
- Distributing Applications and Servers
- Server Integrity during Distribution
- Distribution Descriptor Variables
- Using Distribution in the Ripper Application

Using IcePatch2 to Distribute Applications

In the section so far, "deployment" has meant the creation of descriptors in the registry. A broader definition involves a number of other tasks:

- Writing IceGrid configuration files and preparing data directories on each computer
- Installing the IceGrid binaries and dependent libraries on each computer
- Starting the registry and/or node on each computer, and possibly configuring the systems to launch them automatically
- Distributing your server executables, dependent libraries and supporting files to the appropriate nodes.

The first three tasks are the responsibility of the system administrator, but IceGrid can help with the fourth. Using an IcePatch2 server, you can configure the nodes to download servers automatically and patch them at any time. The illustration below shows the interactions of the components:

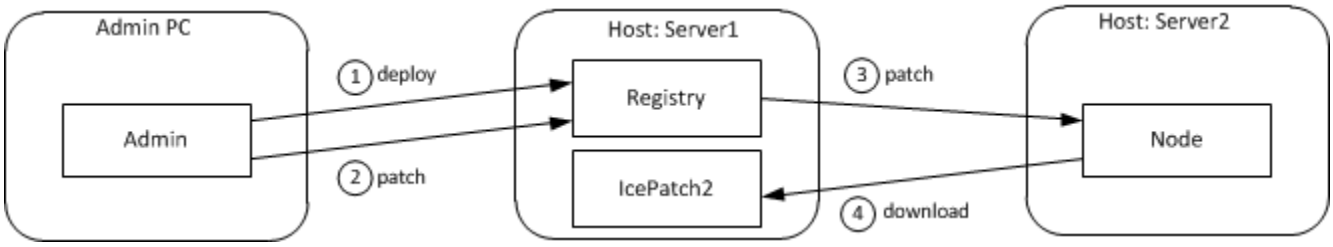

Overview of application distribution.

As you can see, deploying an IceGrid application has greater significance when IcePatch2 is also involved. After deployment, the administrative tool initiates a patch, causing the registry to notify all active nodes that are configured for application distribution to begin the patching process. Since each IceGrid node is an IcePatch2 client, the node performs the patch just like any IcePatch2 client: it downloads everything if no local copy of the distribution exists, otherwise it does an incremental patch in which it downloads only new files and those whose signatures have changed.

The benefits of this feature are clear:

- The distribution files are maintained in a central location
- Updating a distribution on all of the nodes is a simple matter of preparing the master distribution and letting IceGrid do the rest
- Manually transferring executables and supporting files to each computer is avoided, along with the mistakes that manual intervention sometimes introduces.

Deploying an IcePatch2 Server

If you plan to use IceGrid's distribution capabilities, we generally recommend deploying an IcePatch2 server along with your application. Doing so gives you the same benefits as any other IceGrid server, including on-demand activation and remote administration. We will only use one server in our sample application, but you might consider replicating a number of IcePatch2 servers in order to balance the patching load for large distributions.

Patching Considerations

Deploying an IcePatch2 server with your application presents a chicken-and-egg dilemma: how do the nodes download their distributions if the IcePatch2 server is included in the deployment? To answer this question, we need to learn more about IceGrid's behavior.

Deploying and patching are treated as two separate steps: first you deploy the application, then you initiate the patching process. The `icegridadmin` utility combines these steps into one command (`application add`), but also provides an option to disable the patching step if so desired.

Copyright 2016, ZeroC, Inc.

Let's consider the state of the application after deployment but before patching: we have described the servers that run on each node, including file system-dependent attributes such as the pathnames of their executables and default working directories. If these pathnames refer to directories in the distribution, and the distribution has not yet been downloaded to that node, then clearly we cannot attempt to use those servers until patching has completed. Similarly, we cannot deploy an IcePatch2 server whose executable resides in the distribution to be downloaded.

We are ignoring the case where a temporary IcePatch2 server is used to bootstrap other IcePatch2 servers.

For these reasons, we assume that the IcePatch2 server and supporting libraries are distributed by the system administrator along with the IceGrid registry and nodes to the appropriate computers. The server should be configured for on-demand activation so that its node starts it automatically when patching begins. If the server is configured for manual activation, you must start it prior to patching.

IcePatch2 Server Template

The Ice distribution includes an IcePatch2 server template that simplifies the inclusion of IcePatch2 in your application. The relevant portion from the file `config/templates.xml` is shown below:

XML

```
<server-template id="IcePatch2">
    <parameter name="instance-name" default="${application}.IcePatch2"/>
    <parameter name="endpoints" default="default"/>
    <parameter name="directory"/>

    <server id="${instance-name}" exe="icepatch2server"
        application-distrib="false" activation="on-demand">
        <adapter name="IcePatch2" endpoints="${endpoints}">
            <object identity="${instance-name}/server"
type="::IcePatch2::FileServer"/>
        </adapter>
        <adapter name="IcePatch2.Admin" id="" endpoints="tcp -h 127.0.0.1"/>
        <property name="IcePatch2.InstanceName" value="${instance-name}"/>
        <property name="IcePatch2.Directory" value="${directory}"/>
    </server>
</server-template>
```

Notice that the server's pathname is `icepatch2server`, meaning the program must be present in the node's executable search path. The only mandatory parameter is `directory`, which specifies the server's data directory and becomes the value of the `IcePatch2.Directory` property. The value of the `instance-name` parameter is used as the server's identifier when the template is instantiated; its default value includes the name of the application in which the template is used. This identifier also affects the identities of the two well-known objects declared by the server.

Consider the following sample application:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="PatchDemo">
        <node name="Node">
            <server-instance template="IcePatch2" directory="/opt/icepatch2/data"/>
            ...
        </node>
    </application>
</icegrid>
```

Instantiating the `IcePatch2` template creates a server identified as `PatchDemo.IcePatch2` (as determined by the default value for the `instance-name` parameter). The well-known objects use this value as the category in their identities, such as `PatchDemo.IcePatch2/server`.

In order to refer to the `IcePatch2` template in your application, you must have already configured the registry to use the `config/templates.xml` file as your default templates, or copied the template into the XML file describing your application.

Adding Distribution to a Deployment

A distribution descriptor provides the details that a node requires in order to download the necessary files. Specifically, the descriptor supplies the proxy of the IcePatch2 server and the names of the subdirectories comprising the distribution, all of which are optional. If the descriptor does not define the proxy, the following default value is used instead:

```
${application}.IcePatch2/server
```

You may recall that this value matches the default identity configured by the IcePatch2 server template described above. Also notice that this is an indirect proxy, implying that the IcePatch2 server was deployed with the application and can be started on-demand if necessary.

If the descriptor does not select any subdirectories, the node downloads the entire contents of the IcePatch2 data directory.

In XML, a descriptor having the default behavior as described above can be written as shown below:

XML

```
<distrib/>
```

To specify a proxy, use the `icepatch` attribute:

XML

```
<distrib icepatch="PatchDemo.IcePatch2/server"/>
```

Finally, select subdirectories using a nested element:

Copyright 2016, ZeroC, Inc.

XML

```
<distrib>
    <directory>dir1</directory>
    <directory>dir2/subdir</directory>
</distrib>
```

By including only certain subdirectories in a distribution, you are minimizing the time and effort required to download and patch each node. For example, each node in a heterogeneous network might download a platform-specific subdirectory and another subdirectory containing files common to all platforms.

Distributing Applications and Servers

A distribution descriptor can be used in two contexts: within an application, and within a server. When the descriptor appears at the application level, it means every node in the application downloads that distribution. This is useful for distributing files required by all of the nodes on which servers are deployed, especially in a grid of homogeneous computers where it would be tedious to repeat the same distribution information in each server descriptor. Here is a simple XML example:

XML

```
<icegrid>
    <application name="PatchDemo">
        <distrib>
            <directory>Common</directory>
        </distrib>
        ...
    </application>
</icegrid>
```

At the server level, a distribution descriptor downloads the specified directories for the private use of the server:

XML

```
<icegrid>
    <application name="PatchDemo">
        <distrib>
            <directory>Common</directory>
        </distrib>
        <node name="Node">
            <server id="SimpleServer" ...>
                <distrib>
                    <directory>ServerFiles</directory>
                </distrib>
            </server>
        </node>
    </application>
</icegrid>
```


Copyright 2016, ZeroC, Inc.

When a distribution descriptor is defined at both the application and server levels, as shown in the previous example, IceGrid assumes that a dependency relationship exists between the two unless the server is configured otherwise. IceGrid checks this dependency before patching a server; if the server is dependent on the application's distribution, IceGrid patches the application's distribution first, and then proceeds to patch the server's. You can disable this dependency by modifying the server's descriptor:

XML

```
<icegrid>
    <application name="PatchDemo">
        <distrib>
            <directory>Common</directory>
        </distrib>
        <node name="Node">
            <server id="SimpleServer" application-distrib="false" ...>
                <distrib>
                    <directory>ServerFiles</directory>
                </distrib>
            </server>
        </node>
    </application>
</icegrid>
```

Setting the `application-distrib` attribute to `false` informs IceGrid to consider the two distributions independent of one another.

Server Integrity during Distribution

Before an IceGrid node begins patching a distribution, it ensures that all relevant servers are shut down and prevents them from reactivating until patching completes. For example, the node disables all of the servers whose descriptors declare a dependency on the application distribution.

Distribution Descriptor Variables

The node stores application and server distributions in its data directory. The path names of the distributions are represented by reserved variables that you can use in your descriptors:

- `application.distrib`
  This variable can be used within server descriptors to refer to the top-level directory of the application distribution.
- `server.distrib`
  The value of this variable is the top-level directory of a server distribution. It can be used only within a server descriptor that has a distribution.

The XML example shown below illustrates the use of these variables:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="PatchDemo">
        <distrib>
            <directory>Common</directory>
        </distrib>
        <node name="Node">
            <server id="Server1" exe="${application.distrib}/Common/Bin/Server1"
...>
            </server>
            <server id="Server2" exe="${server.distrib}/Server2Files/Bin/Server2"
...>
                <option>-d</option>
                <option>${server.distrib}/Server2Files</option>
                <distrib>
                    <directory>Server2Files</directory>
                </distrib>
            </server>
        </node>
    </application>
</icegrid>
```

Notice that the descriptor for `Server2` supplies the server's distribution directory as command-line options.

Using Distribution in the Ripper Application

Adding an application distribution to our ripper example requires two minor changes to our descriptors:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="Ripper">
        <replica-group id="EncoderAdapters">
            <load-balancing type="adaptive"/>
            <object identity="EncoderFactory" type="::Ripper::MP3EncoderFactory"/>
        </replica-group>
        <server-template id="EncoderServerTemplate">
            <parameter name="index"/>
            <parameter name="exepath" default="/opt/ripper/bin/server"/>
            <server id="EncoderServer${index}" exe="${exepath}"
activation="on-demand">
                <adapter name="EncoderAdapter" replica-group="EncoderAdapters"
                    endpoints="tcp"/>
            </server>
        </server-template>
        <distrib/>
        <node name="Node1">
            <server-instance template="EncoderServerTemplate" index="1"/>
            <server-instance template="IcePatch2"
directory="/opt/ripper/icepatch"/>
        </node>
        <node name="Node2">
            <server-instance template="EncoderServerTemplate" index="2"/>
        </node>
    </application>
</icegrid>
```

An application distribution is sufficient for this example because we are deploying the same server on each node. We have also deployed an IcePatch2 server on `Node1` using the template.

See Also

- IcePatch2
- icegridadmin Command Line Tool
- IceGrid Templates
- Distrib Descriptor Element
- Using Descriptor Variables and Parameters
- Object Adapter Replication
- IcePatch2 Object Identities

Copyright 2016, ZeroC, Inc.

IceGrid Administrative Sessions

To access IceGrid's administrative facilities from a program, you must first establish an administrative session. Once done, a wide range of services are at your disposal, including the manipulation of IceGrid registries, nodes, and servers; deployment of new components such as well-known objects; and dynamic monitoring of IceGrid events.

Note that, for replicated registries, an administrative session can be established with either the master or a slave registry replica, but a session with a slave replica is restricted to read-only operations.

On this page:

- Creating an Administrative Session
- Accessing Log Files Remotely
- Dynamic Monitoring in IceGrid
  - Observer Interfaces
  - Registering Observers

Creating an Administrative Session

The `Registry` interface provides two operations for creating an administrative session:

Slice

```
module IceGrid {
    exception PermissionDeniedException {
        string reason;
    };

    interface Registry {
        AdminSession* createAdminSession(string userId, string password)
            throws PermissionDeniedException;

        AdminSession* createAdminSessionFromSecureConnection()
            throws PermissionDeniedException;

        idempotent int getSessionTimeout();

        idempotent int getACMTimeout();
        // ...
    };
};
```

The `createAdminSession` operation expects a username and password and returns a session proxy if the client is allowed to create a session. By default, IceGrid does not allow the creation of administrative sessions. You must define the property `IceGrid.Registry.AdminPermissionsVerifier` with the proxy of a permissions verifier object to enable session creation with `createAdminSession`. The verifier object must implement the interface `Glacier2::PermissionsVerifier`.

The `createAdminSessionFromSecureConnection` operation does not require a username and password because it uses the credentials supplied by an SSL connection to authenticate the client. As with `createAdminSession`, you must configure the proxy of a permissions verifier object before clients can use `createAdminSessionFromSecureConnection` to create a session. In this case, the `IceGrid.Registry.AdminSSLPermissionsVerifier` property specifies the proxy of a verifier object that implements the interface `Glacier2::SSLPermissionsVerifier`.

As an example, the following code demonstrates how to obtain a proxy for the registry and invoke `createAdminSession`:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx base = communicator->stringToProxy("IceGrid/Registry");
IceGrid::RegistryPrx registry = IceGrid::RegistryPrx::checkedCast(base);
string username = ...;
string password = ...;
IceGrid::AdminSessionPrx session;
try {
    session = registry->createAdminSession(username, password);
} catch (const IceGrid::PermissionDeniedException & ex) {
    cout << "permission denied:\n" << ex.reason << endl;
}
```

The `AdminSession` interface provides operations for accessing log files and establishing observers. Furthermore, two additional operations are worthy of your attention:

Slice

```
module IceGrid {
    interface AdminSession extends Glacier2::Session {
        idempotent void keepAlive();
        idempotent Admin* getAdmin();
        // ...
    };
};
```

If your program uses an administrative session indefinitely, you must prevent the session from expiring. You have two options for keeping a session alive:

1. Call `Registry::getSessionTimeout` and periodically invoke `AdminSession::keepAlive` every *timeout* (or less) seconds
2. Call `Registry::getACMTimeout` and configure ACM settings on the connection

We recommend using the second approach, which you can implement as follows:

C++

```
int acmTimeout = registry->getACMTimeout();
if (acmTimeout > 0) {
    Ice::ConnectionPtr conn = session->ice_getCachedConnection();
    conn->setACM(acmTimeout, IceUtil::None, Ice::HeartbeatAlways);
}
```

Enabling heartbeats on the connection causes Ice to automatically send a heartbeat message at regular intervals determined by the given timeout value. The server ignores these messages, but they serve the purpose of keeping the session alive.

The ACM features shown here were introduced in Ice 3.6.

The `getAdmin` operation returns a proxy for the `IceGrid::Admin` interface, which provides complete access to the registry's settings. For

Copyright 2016, ZeroC, Inc.

this reason, you must use extreme caution when enabling administrative sessions.

Accessing Log Files Remotely

IceGrid's `AdminSession` interface provides operations for remotely accessing the log files of a registry, node, or server:

Slice

```
module IceGrid {
interface AdminSession extends Glacier2::Session {
    // ...

    FileIterator* openServerLog(string id, string path, int count)
        throws FileNotAvailableException, ServerNotExistException,
               NodeUnreachableException, DeploymentException;
    FileIterator* openServerStdErr(string id, int count)
        throws FileNotAvailableException, ServerNotExistException,
               NodeUnreachableException, DeploymentException;
    FileIterator* openServerStdOut(string id, int count)
        throws FileNotAvailableException, ServerNotExistException,
               NodeUnreachableException, DeploymentException;

    FileIterator* openNodeStdErr(string name, int count)
        throws FileNotAvailableException, NodeNotExistException,
               NodeUnreachableException;
    FileIterator* openNodeStdOut(string name, int count)
        throws FileNotAvailableException, NodeNotExistException,
               NodeUnreachableException;

    FileIterator* openRegistryStdErr(string name, int count)
        throws FileNotAvailableException,
               RegistryNotExistException,
               RegistryUnreachableException;
    FileIterator * openRegistryStdOut(string name, int count)
        throws FileNotAvailableException,
               RegistryNotExistException,
               RegistryUnreachableException;
};
};
```

In order to access the text of a program's standard output or standard error log, you must configure it using the `Ice.StdOut` and `Ice.StdErr` properties, respectively. For registries and nodes, you must define these properties explicitly but, for servers, the node defines these properties automatically if the property `IceGrid.Node.Output` is defined, causing the server's output to be logged in individual files.

If `IceGrid.Node.Output` is *not* defined, the following rules apply:

- If the node is started from a console or shell, servers share the node's `stdout` and `stderr`. If `Ice.StdOut` and/or `Ice.StdErr` properties are defined for the node, the servers' output is redirected to the specified files as well.
- If the node is started as a Unix daemon and `--noclose` is not used, the servers' output is lost, except if `Ice.StdOut` and/or `Ice.StdErr` properties are set for the node, in which case the servers' output is redirected to the specified files.
- If the node is started as a Windows service, the servers' output is lost even if `Ice.StdOut` and/or `Ice.StdErr` are set.

Log messages from the node itself are sent to `stderr` unless you set `Ice.UseSyslog` (for Unix). If the node is started as a Windows service, its log messages always are sent to the Windows event log.

In the case of `openServerLog`, the value of the `path` argument must resolve to the same file as one of the server's log descriptors. This security measure prevents a client from opening an arbitrary file on the server's host.

Copyright 2016, ZeroC, Inc.

All of the operations accept a `count` argument and return a proxy to a `FileIterator` object. The `count` argument determines where to start reading the log file: if the value is negative, the iterator is positioned at the beginning of the file, otherwise the iterator is positioned to return the last `count` lines of text.

The `FileIterator` interface is quite simple:

Slice

```
module IceGrid {
interface FileIterator {
    bool read(int size, out Ice::StringSeq lines)
        throws FileNotAvailableException;
    void destroy();
};
};
```

A client may invoke the `read` operation as many times as necessary. The `size` argument specifies the maximum number of bytes that `read` can return; the client must not use a size that would cause the reply to exceed the client's configured maximum message size.

If this is the client's first call to `read`, the `lines` argument holds whatever text was available from the iterator's initial position, and the iterator is repositioned in preparation for the next call to `read`. The operation returns false to indicate that more text is available and true if all available text has been read.

Line termination characters are removed from the contents of `lines`. When displaying the text, you must be aware that the first and last elements of the sequence can be partial lines. For example, the last line of the sequence might be incomplete if the limit specified by `size` is reached. The next call to `read` returns the remainder of that line as the first element in the sequence.

As an example, the C++ code below displays the contents of a log file and waits for new text to become available:

C++

```
IceGrid::FileIteratorPrx iter = ...;
while(true) {
    Ice::StringSeq lines;
    bool end = iter->read(10000, lines);
    if (!lines.empty()) {
        // The first line might be a continuation from
        // the previous call to read.
        cout << lines[0];
        for (Ice::StringSeq::const_iterator p =
++lines.begin(); p != lines.end(); ++p)
            cout << endl << *p << flush;
    }
    if (end)
        sleep(1);
}
```

Notice that the loop includes a delay in case `read` returns true, which prevents the client from entering a busy loop when no data is currently available.

The client should call `destroy` when the iterator object is no longer required. At the time the client's session terminates, IceGrid reclaims any iterators that were not explicitly destroyed.

If the client waits for new data, it must take steps to prevent the administrative session from expiring.

Copyright 2016, ZeroC, Inc.

With these operations, an administrative client can retrieve any text file on a system where an IceGrid node is running. While it's common for this text file to contain the output of an Ice logger, it could contain other unrelated outputs. This text file is presented as a sequence of strings, with no particular structure for these strings.

The `Logger` admin facet provides another, and often better, way to retrieve the log messages sent to the logger of a server, node or registry. With the `Logger` facet, you retrieve log messages–and only log messages–as typed structures, whether or not the logger's output is stored in a text file or stored at all. With the `AdminSession` file operations presented above, you can retrieve the log files of a server even when this server is not running; conversely, with the `Logger` admin facet, the target server must be running since this `Logger` facet is hosted in that server. Finally, the programming style espoused by the `AdminSession` file operations is a pull model: the administrative client calls `read` from time to time on the `FileIterator` object provided by IceGrid. With the `Logger` admin facet, the administrative client uses a push model: it registers a remote logger object with the target server, node and registry, and this remote logger receives new log messages as soon as they are generated.

Dynamic Monitoring in IceGrid

IceGrid allows an application to monitor relevant state changes by registering callback objects. (The IceGrid GUI tool uses these callback interfaces for its implementation.) The callback interfaces are useful to, for example, automatically generate an email notification when a node goes down or some other state change of interest occurs.

Observer Interfaces

IceGrid offers a callback interface for each major component of the IceGrid architecture:

Copyright 2016, ZeroC, Inc.

Slice

```
module IceGrid {
interface NodeObserver {
    void nodeInit(NodeDynamicInfoSeq nodes);
    void nodeUp(NodeDynamicInfo node);
    void nodeDown(string name);
    void updateServer(string node, ServerDynamicInfo updatedInfo);
    void updateAdapter(string node, AdapterDynamicInfo updatedInfo);
};

interface ApplicationObserver {
    void applicationInit(int serial, ApplicationInfoSeq applications);
    void applicationAdded(int serial, ApplicationInfo desc);
    void applicationRemoved(int serial, string name);
    void applicationUpdated(int serial, ApplicationUpdateInfo desc);
};

interface AdapterObserver {
    void adapterInit(AdapterInfoSeq adpts);
    void adapterAdded(AdapterInfo info);
    void adapterUpdated(AdapterInfo info);
    void adapterRemoved(string id);
};

interface ObjectObserver {
    void objectInit(ObjectInfoSeq objects);
    void objectAdded(ObjectInfo info);
    void objectUpdated(ObjectInfo info);
    void objectRemoved(Ice::Identity id);
};

interface RegistryObserver {
    void registryInit(RegistryInfoSeq registries);
    void registryUp(RegistryInfo node);
    void registryDown(string name);
};
};
```

The next section describes how to install an observer.

Registering Observers

The `AdminSession` interface provides two operations for registering your observers:

Copyright 2016, ZeroC, Inc.

Slice

```
module IceGrid {
    interface AdminSession extends Glacier2::Session {
        idempotent void keepAlive();

        idempotent void setObservers(
                RegistryObserver* registryObs,
                NodeObserver* nodeObs,
                ApplicationObserver* appObs,
                AdapterObserver* adptObs,
                ObjectObserver* objObs)
            throws ObserverAlreadyRegisteredException;

        idempotent void setObserversByIdentity(
                Ice::Identity registryObs,
                Ice::Identity nodeObs,
                Ice::Identity appObs,
                Ice::Identity adptObs,
                Ice::Identity objObs)
            throws ObserverAlreadyRegisteredException;

        // ...
    };
};
```

You should invoke `setObservers` and supply proxies when it is possible for the registry to establish a separate connection to the client to deliver its callbacks. If network restrictions such as firewalls prevent such a connection, you should use the `setObserversByIdentity` operation, which creates a bidirectional connection instead.

You can pass a null proxy for any parameter to `setObservers`, or an empty identity for any parameter to `setObserversByIdentity`, if you want to use only some of the observers. In addition, passing a null proxy or an empty identity for an observer cancels a previous registration of that observer. The operations raise `ObserverAlreadyRegisteredException` if you pass a proxy or identity that was registered in a previous call.

Once the observers are registered, operations corresponding to state changes will be invoked on the observers. (See the Slice API Reference for details on the data passed to the observers. You can also look at the source code for the IceGrid GUI implementation in the Ice for Java distribution to see how observers are used by the GUI.)

Finally, remember to take the steps necessary to prevent the administrative session from expiring.

See Also

- Registry Replication
- Securing a Glacier2 Router
- Resource Allocation using IceGrid Sessions
- Log Descriptor Element
- icegridadmin Command Line Tool
- Bidirectional Connections
- IceGrid.*
- Logger admin Facet

Copyright 2016, ZeroC, Inc.

Glacier2 Integration with IceGrid

This section provides information on integrating a Glacier2 router into your IceGrid environment.

On this page:

- Configuration Changes for using Glacier2 with IceGrid
- Remote IceGrid Administration via Glacier2
- Resource Allocation using Glacier2 and IceGrid
- Session Considerations for Glacier2 and IceGrid
- Deploying Glacier2 with IceGrid

Configuration Changes for using Glacier2 with IceGrid

A typical IceGrid client must be configured with a locator proxy, but the configuration requirements change when the client accesses the location service indirectly via a Glacier2 router as shown below:

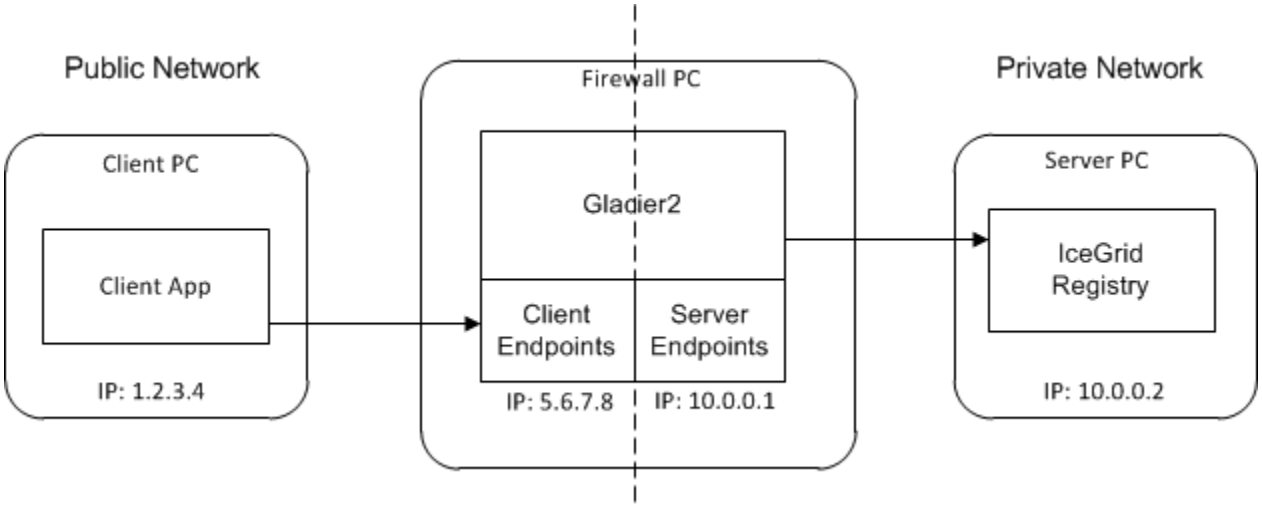

Using IceGrid via a Glacier2 router.

In this situation, it is the router that must be configured with a locator proxy.

Assuming the registry's client endpoint in the illustration uses port `8000`, the router requires the following setting for the `Ice.Default.Locator` property:

```
Ice.Default.Locator=IceGrid/Locator:tcp -h 10.0.0.2 -p 8000
```

Fortunately, the node supplies this property when it starts the router, so there is no need to configure it explicitly. Note that all of the router's clients use the same locator.

Remote IceGrid Administration via Glacier2

If you intend to administer IceGrid remotely via a Glacier2 router, you must define one of the following properties (or both), depending on whether you use user name and password authentication or a secure connection:

```
Glacier2.SessionManager=IceGrid/AdminSessionManager
Glacier2.SSLSessionManager=IceGrid/AdminSSLSessionManager
```

These session managers are accessible via the registry's administrative session manager endpoints, so the Glacier2 router must be authorized to establish a connection to these endpoints. Note that you must secure these endpoints, otherwise arbitrary clients can manipulate the session managers. An administrative session is allowed to access any object by default. To restrict access to the `IceGrid:`

Copyright 2016, ZeroC, Inc.

`:AdminSession` object and the `IceGrid::Admin` object that is returned by the session's `getAdmin` operation, you must set the property `IceGrid.Registry.AdminSessionFilters` to one.

Resource Allocation using Glacier2 and IceGrid

To allocate servers and objects, a program can establish a client session via Glacier2. Depending on the authentication method, one or both of the following properties must be set in the Glacier2 configuration:

```
Glacier2.SessionManager=IceGrid/SessionManager
Glacier2.SSLSessionManager=IceGrid/SSLSessionManager
```

These session managers are accessible via the registry's session manager endpoints, so the Glacier2 router must be authorized to establish a connection to these endpoints.

A client session is allowed to access any object by default. To restrict access to the `IceGrid::Session` and `IceGrid::Query` objects, you must set the property `IceGrid.Registry.SessionFilters` to one. However, you can use the allocation mechanism to access additional objects and adapters. IceGrid adds an identity filter when a client allocates an object and removes that filter again when the object is released. When a client allocates a server, IceGrid adds an adapter identity filter for the server's indirect adapters and removes that filter again when the server is released.

Session Considerations for Glacier2 and IceGrid

Providing access to administrative sessions and client sessions both require that you define at least one of the properties `Glacier2.SessionManager` and `Glacier2.SSLSessionManager`, which presents a potential problem if you intend to access both types of sessions via the same Glacier2 router.

The simplest solution is to dedicate a router instance to each type of session. However, if you need to access both types of sessions from a single router, you can accomplish it only if you use a different authentication mechanism for each type of session. For example, you can configure the router as follows:

```
Glacier2.SessionManager=IceGrid/SessionManager
Glacier2.SSLSessionManager=IceGrid/AdminSSLSessionManager
```

This configuration uses user name and password authentication for client sessions, and SSL authentication for administrative sessions. If this restriction is too limiting, you must use two router instances.

Deploying Glacier2 with IceGrid

The Ice distribution includes default server templates for Ice services such as IcePatch2 and Glacier2 that simplify the task of deploying these servers in an IceGrid domain.

The relevant portion from the file `config/template.xml` is shown below:

Copyright 2016, ZeroC, Inc.

XML

```
<server-template id="Glacier2">
    <parameter name="instance-name" default="${application}.Glacier2"/>
    <parameter name="client-endpoints"/>
    <parameter name="server-endpoints"/>
    <parameter name="session-timeout" default="0"/>

    <server id="${instance-name}" exe="glacier2router">
    <properties>
        <property name="Glacier2.Client.Endpoints" value="${client-endpoints}"/>
        <property name="Glacier2.Server.Endpoints" value="${server-endpoints}"/>
        <property name="Glacier2.InstanceName" value="${instance-name}"/>
        <property name="Glacier2.SessionTimeout" value="${session-timeout}"/>
    </properties>
</server-template>
```

Notice that the server's pathname is `glacier2router`, meaning the program must be present in the node's executable search path. Another important point is the server's activation mode: it uses manual activation (the default), meaning the router must be started manually. This requirement becomes clear when you consider that the router is the point of contact for remote clients; if the router is not running, there is no way for a client to contact the locator and cause the router to be started on-demand.

The template defines only a few properties; if you want to set additional properties, you can define them in the server instance property set.

Of interest is the `instance-name` parameter, which allows you to configure the `Glacier2.InstanceName` property. The parameter's default value includes the name of the application in which the template is used. This parameter also affects the identities of the objects implemented by the router.

Consider the following sample application:

```
<icegrid>
    <application name="Glacier2Demo">
        <node name="Node">
            <server-instance template="Glacier2"
                client-endpoints="tcp -h 5.6.7.8 -p 8000"
                session-timeout="300"
                server-endpoints="tcp -h 10.0.0.1"/>
            ...
        </node>
    </application>
</icegrid>
```

Instantiating the `Glacier2` template creates a server identified as `Glacier2Demo.Glacier2` (as determined by the default value for the `instance-name` parameter). The router's objects use this value as the category in their identities, such as `Glacier2Demo.Glacier2/router`. The router proxy used by clients must contain a matching identity.

In order to refer to the `Glacier2` template in your application, you must have already configured the registry to use the `config/templates.xml` file as your default templates, or copied the template into the XML file describing your application.

Note that IceGrid cannot start a Glacier2 router if the router's security configuration requires that a passphrase be entered. In this situation, you have no choice but to start the router yourself so that you can provide the passphrase when prompted.

See Also

Copyright 2016, ZeroC, Inc.

- Glacier2
- IceGrid Templates
- Getting Started with Glacier2
- Glacier2.*
- IceGrid.*

Copyright 2016, ZeroC, Inc.

IceGrid XML Reference

This section provides a reference for the XML elements that define IceGrid descriptors, in alphabetical order.

IceGrid XML files must use UTF-8 encoding.

Topics

- Adapter Descriptor Element
- Allocatable Descriptor Element
- Application Descriptor Element
- DbEnv Descriptor Element
- DbProperty Descriptor Element
- Description Descriptor Element
- Directory Descriptor Element
- Distrib Descriptor Element
- IceBox Descriptor Element
- IceGrid Descriptor Element
- Load-Balancing Descriptor Element
- Log Descriptor Element
- Node Descriptor Element
- Object Descriptor Element
- Parameter Descriptor Element
- Properties Descriptor Element
- Property Descriptor Element
- Replica-Group Descriptor Element
- Server Descriptor Element
- Server-Instance Descriptor Element
- Server-Template Descriptor Element
- Service Descriptor Element
- Service-Instance Descriptor Element
- Service-Template Descriptor Element
- Variable Descriptor Element
- Using Command Line Options in Descriptors
- Setting Environment Variables in Descriptors

See Also

- Using IceGrid Deployment

Copyright 2016, ZeroC, Inc.

Adapter Descriptor Element

An `adapter` element defines an indirect object adapter.

This element may only appear as a child of a `server` or `service` element.

The following attributes are supported:

Attribute	Description	Required
`endpoints`	Specifies one or more endpoints for this object adapter. These endpoints typically do not specify a port. This attribute is translated into a definition of the adapter's `Endpoints` configuration property	No
`id`	Specifies an object adapter identifier. The identifier must be unique among all adapters and replica groups in the registry. This attribute is translated into a definition of the adapter's `AdapterId` configuration property. If not defined, a default value is constructed from the adapter name and server ID (and service name for an IceBox service).	Yes
`name`	The name of the object adapter as used in the server that creates it.	Yes
`priority`	Specifies the priority of the object adapter as an integer value. The object adapter priority is used by the `Ordered` replica group load balancing policy to determine the order of the endpoints returned by a locate request. Endpoints are ordered from the smallest priority value to the highest. If not defined, the value is 0.	No
`proxy-options`	The proxy options to use when generating the adapter well-known or allocatable object proxies. This attribute is also translated into a definition of the adapter's ProxyOptions configuration property.	No
`register-process`	This attribute is only valid if the enclosing server uses an Ice version prior to 3.3. In Ice 3.3 or later, this functionality is replaced by the administrative facility. If the value is `true`, this object adapter registers an object in the IceGrid registry that facilitates graceful shutdown of the server. Only one object adapter in a server should set this attribute to `true`. If not defined, the default value is `false`.	No
`replica-group`	Specifies a replica group identifier. A non-empty value signals that this object adapter is a member of the indicated replica group. This attribute is translated into a definition of the adapter's `ReplicaGroupId` configuration property. If not defined, the default value is an empty string.	No
`server-lifetime`	A value of `true` indicates that the lifetime of this object adapter is the same as the lifetime of its server. This information is used by the IceGrid node to determine the state of the server. Specifically, the server is considered activated when all the object adapters with the `server-lifetime` attribute set to `true` are registered with the registry (the object adapter registers itself during activation). Furthermore, server deactivation is considered to begin as soon as one object adapter with the `server-lifetime` attribute set to `true` is unregistered with the registry (the object adapter unregisters itself during deactivation).	No

An optional nested `description` element provides free-form descriptive text.

Here is an example to demonstrate the use of this element.

XML

```
<adapter name="MyAdapter"
    endpoints="default"
    id="MyAdapterId"
    proxy-options="-t -e 1.0"
    replica-group="MyReplicaGroup">
    <description>A description of the adapter.</description>
    ...
</adapter>
```

See Also

Copyright 2016, ZeroC, Inc.

- Object Adapters
- Server Descriptor Element
- Service Descriptor Element
- Object Adapter Endpoints
- Object Adapter Replication
- Creating an Object Adapter
- Load Balancing
- IceGrid Server Activation
- Object Adapter Properties

Copyright 2016, ZeroC, Inc.

Allocatable Descriptor Element

An allocatable element defines an allocatable object in the IceGrid registry. Clients can allocate this object using only its identity, or they can allocate objects of a specific type.

This element may only appear as a child of an `adapter` element.

The following attributes are supported:

Attribute	Description	Required
`identity`	Specifies the identity by which this allocatable object is known.	Yes
`property`	Specifies the name of a property to generate that contains the stringified identity of this allocatable.	No
`type`	An arbitrary string used to group allocatable objects. By convention, the string represents the most-derived Slice type ID of the object, but an application is free to use another convention.	No
`proxy-options`	The proxy options to use for the proxy of the allocatable object returned by IceGrid.	No

See Also

- Resource Allocation using IceGrid Sessions
- Adapter Descriptor Element
- Type IDs

Copyright 2016, ZeroC, Inc.

Application Descriptor Element

An `application` element defines an IceGrid application. An application typically contains at least one `node` element, but it may also be used for other purposes such as defining server and service templates, default templates, replica groups and property sets.

This element must be a child of an `icegrid` element. Only one `application` element is permitted per file.

The following attributes are supported.

Attribute	Description	Required
`import-default-templates`	If `true`, the default templates configured for the IceGrid registry are imported and available for use within this application. If not specified, the default value is `false`.	No
`name`	The name of the application. This name must be unique among all applications in the registry. Within the application, child elements can refer to its name using the reserved variable `${application}`.	Yes

An optional nested `description` element provides free-form descriptive text.

Here is an example to demonstrate the use of this element:

XML

```
<icegrid>
    <application name="MyApplication" import-default-templates="true">
        <description>A description of the application.</description>
        ...
    </application>
</icegrid>
```

See Also

- IceGrid Architecture
- Node Descriptor Element
- Server-Template Descriptor Element
- Service-Template Descriptor Element
- Replica-Group Descriptor Element
- Properties Descriptor Element
- IceGrid Descriptor Element
- Description Descriptor Element
- IceGrid Templates
- Using Descriptor Variables and Parameters

Copyright 2016, ZeroC, Inc.

DbEnv Descriptor Element

A `dbenv` element causes an IceGrid node to generate Freeze configuration properties for the server or service in which it is defined. The IceGrid node also creates a directory for the Berkeley DB database environment if needed.

This element may contain `dbproperty` elements.This element may only appear as a child of a `server` or `service` element.

The following attributes are supported:

Attribute	Description	Required
`home`	Specifies the directory to use as the database environment. If not defined, a sub-directory in the node's data directory is used.	No
`name`	The name of the Berkeley DB database environment.	Yes

The values of the `name` and `home` attributes are used to define the `Freeze.DbEnv.`*`env-name`*`.DbHome` property as shown below:

`Freeze.DbEnv.`*`name`*`.DbHome=`*`home`*

An optional nested `description` element provides free-form descriptive text.

Here is an example to demonstrate the use of this element:

XML

```
<dbenv name="MyEnv" home="/opt/data/db">
    <description>A description of the Freeze/BDB
database environment.</description>
    ...
</dbenv>
```

See Also

- Freeze
- DbProperty Descriptor Element
- Server Descriptor Element
- Service Descriptor Element
- Description Descriptor Element
- Freeze.*

Copyright 2016, ZeroC, Inc.

DbProperty Descriptor Element

A `dbproperty` element defines a Berkeley DB configuration property for a Freeze database environment.

This element may only appear as a child of a `dbenv` element.

The following attributes are supported:

Attribute	Description	Required
`name`	The name of the configuration property.	Yes
`value`	The value of the configuration property. If not defined, the value is an empty string.	No

Here is an example to demonstrate the use of this element:

XML

```
<dbenv name="MyEnv" home="/opt/data/db">
    <description>A description of the database environment.</description>
    <dbproperty name="set_cachesize" value="0 52428800 1"/>
</dbenv>
```

See Also

- Freeze
- DbEnv Descriptor Element

Copyright 2016, ZeroC, Inc.

Description Descriptor Element

A `description` element specifies a description of its parent element.

This element may only appear as a child of the `application`, `replica-group`, `node`, `server`, `service`, `icebox`, `adapter`, and `dbenv` elements.

Here is an example to demonstrate the use of this element:

XML

```
<node name="localnode">
    <description>Free form descriptive text for localnode</description>
</node>
```

See Also

- Application Descriptor Element
- Replica-Group Descriptor Element
- Node Descriptor Element
- Server Descriptor Element
- Service Descriptor Element
- IceBox Descriptor Element
- Adapter Descriptor Element
- DbEnv Descriptor Element

Copyright 2016, ZeroC, Inc.

Directory Descriptor Element

A `directory` element specifies a directory in a distribution.

This element may only appear as a child of the `distrib` element.

This element supports no attributes.
See Also

- Application Distribution
- Distrib Descriptor Element

Copyright 2016, ZeroC, Inc.

Distrib Descriptor Element

A `distrib` element specifies the distribution files to download from an IcePatch2 server as well as the server's proxy.

This element may only appear as a child of an `application` element or a `server` element.

The following attributes are supported:

Attribute	Description	Required
`icepatch`	Proxy for the IcePatch2 server. If not defined, the default value is `${application}.IcePatch2/server`.	No

Here is an example to demonstrate the use of this element:

XML

```
<distrib icepatch="DemoIcePatch2/server:tcp -p 12345">
    <directory>dir1</directory>
    <directory>dir2/subdir</directory>
</distrib>
```

See Also

- Application Distribution
- IcePatch2
- Application Descriptor Element
- Server Descriptor Element

Copyright 2016, ZeroC, Inc.

IceBox Descriptor Element

An `icebox` element defines an IceBox server to be deployed on a node. It typically contains at least one `service` element, and may supply additional information such as command-line options, environment variables, configuration properties and a server distribution.

This element may only appear as a child of a `node` element or a `server-template` element.

This element supports the same attributes as the `server` element.

An optional nested `description` element provides free-form descriptive text.

Here is an example to demonstrate the use of this element:

XML

```
<icebox id="MyIceBox"
        activation="on-demand"
        activation-timeout="60"
        application-distrib="false"
        deactivation-timeout="60"
        exe="/opt/Ice/bin/icebox"
        pwd="/">
    <option>--Ice.Trace.Network=1</option>
    <env>PATH=/opt/Ice/bin:$PATH</env>
    <property name="IceBox.UseSharedCommunicator.Service1" value="1"/>
    <service name="Service1" .../>
    <service-instance template="ServiceTemplate" name="Service2"/>
</icebox>
```

See Also

- IceBox
- Service Descriptor Element
- Adapter Descriptor Element
- Properties Descriptor Element
- Node Descriptor Element
- Server-Template Descriptor Element
- Server Descriptor Element
- Description Descriptor Element
- Using Command Line Options in Descriptors
- Setting Environment Variables in Descriptors
- Application Distribution
- Administrative Facility

Copyright 2016, ZeroC, Inc.

IceGrid Descriptor Element

The `icegrid` element is the top-level element for IceGrid descriptors in XML files. This elements supports no attributes.
See Also

- Using IceGrid Deployment

Copyright 2016, ZeroC, Inc.

Load-Balancing Descriptor Element

A `load-balancing` element determines the load balancing policy used by a replica group.

This element may only appear as a child of a `replica-group` element.

The following attributes are supported:

Attribute	Description	Required
`load-sample`	Specifies the load sample to use for the adaptive load balancing policy. It can only be defined if `type` is set to `adaptive`. Legal values are `1`, `5` and `15`. If not specified, the load sample default value is `1`.	No
`n-replicas`	Specifies the maximum number of replicas used to compute the endpoints of the replica group. If not specified, the default value is `1`.	No
`type`	Specifies the type of load balancing. Legal values are `adaptive`, `ordered`, `round-robin` and `random`.	Yes

Here is an example to demonstrate the use of this element:

XML

```
<application name="MyApp">
    <replica-group id="ReplicatedAdapter">
        <load-balancing type="adaptive" load-sample="15" n-replicas="3"/>
        <description>A description of this replica group.</description>
        <object identity="WellKnownObject" .../>
    </replica-group>
    ...
</application>
```

See Also

- Load Balancing
- Object Adapter Replication
- Replica-Group Descriptor Element

Copyright 2016, ZeroC, Inc.

Log Descriptor Element

A `log` element specifies the name of a log file for a server or service. A `log` element must be defined for each log file that can be accessed remotely by an administrative tool. Note that it is not necessary to define a `log` element for the values of the `Ice.StdErr` and `Ice.StdOut` properties.

This element may only appear as a child of a `server` element or a `service` element.

The following attributes are supported:

Attribute	Description	Required
`path`	The path name of the log file. If a relative path is specified, it is relative to the current working directory of the node. The node must have sufficient access privileges to read the file.	Yes
`property`	Specifies the name of a property in which to store the path name of the log file as given in the `path` attribute.	No

Here is an example to demonstrate the use of this element:

XML

```
<server id="MyServer" ...>
    <log path="${server}.log" property="LogFile"/>
</server>
```

See Also

- IceGrid Administrative Sessions
- Server Descriptor Element
- Service Descriptor Element

Copyright 2016, ZeroC, Inc.

Node Descriptor Element

A `node` element defines an IceGrid node. The servers that the node is responsible for managing are described in child elements.

This element may only appear as a child of an `application` element. Multiple `node` elements having the same name may appear in an application. Their contents are merged and the last definition of `load-factor` has precedence.

The following attributes are supported:

Attribute	Description	Required
`load-factor`	A floating point value defining the factor that is multiplied with the node's load average. The load average is used by the adaptive load balancing policy to figure out which node is the least loaded. The default is `1.0` on Unix platforms and `1/NCPUS` on Windows (where `NCPUS` is the number of CPUs in the node's computer). Note that, if Unix and Windows machines are part of a replica group, the Unix and Windows figures are not directly comparable, but the registry still makes an attempt to pick the least-loaded node.	No
`name`	The name must be unique among all nodes in the registry. Within the node, child elements can refer to its name using the reserved variable `${node}`. An `icegridnode` process representing this node must be started on the desired computer and its configuration property `IceGrid.Node.Name` must match this attribute.	Yes

Here is an example to demonstrate the use of this element:

XML

```
<node name="Node1" load-factor="2.0">
    <description>A description of this node.</description>
    <server id="Server1" ...>
        <property name="NodeName" value="${node}"/>
        ...
    </server>
</node>
```

See Also

- Application Descriptor Element
- Load Balancing
- Object Adapter Replication
- icegridnode

Copyright 2016, ZeroC, Inc.

Object Descriptor Element

An `object` element defines a well-known object in the IceGrid registry. Clients can refer to this object using only its identity, or they can search for well-known objects of a specific type.

This element may only appear as a child of an `adapter` element or a `replica-group` element.

The following attributes are supported:

Attribute	Description	Required
`identity`	Specifies the identity by which this object is known.	Yes
`property`	Specifies the name of a property to generate that contains the stringified identity of the object. This attribute is only allowed if this `object` element is a child of an `adapter` element.	No
`type`	An arbitrary string used to group objects. By convention, the string represents the most-derived Slice type ID of the object, but an application is free to use another convention.	No
`proxy-options`	The proxy options to use for the proxy of the well-known object return by IceGrid.	No

Here is an example to demonstrate the use of this element:

XML

```
<adapter name="MyAdapter" id="WellKnownAdapter" ...>
    <object identity="WellKnownObject" type="::Module::WellKnownInterface"
proxy-options="-o"/>
</adapter>
```

In the configuration above, the object can be located via the equivalent proxies `WellKnownObject` and `WellKnownObject@WellKnownAdapter`.

See Also

- Well-Known Objects
- Adapter Descriptor Element
- Replica-Group Descriptor Element
- Type IDs

Copyright 2016, ZeroC, Inc.

Parameter Descriptor Element

A `parameter` element defines a template parameter. Template parameters must be declared with this element to be used in template instantiation.

This element may only appear as a child of a `server-template` element or a `service-template` element.

The following attributes are supported:

Attribute	Description	Required
`name`	The name of the parameter. For example, if `index` is used as the name of a parameter, it can be referenced using `${index}` in the server or service template.	Yes
`default`	An optional default value for the parameter. This value is used if the parameter is not defined when a server or service is instantiated.	No

Here is an example to demonstrate the use of this element:

XML

```
<server-template id="MyServerTemplate">
    <parameter name="index"/>
    <parameter name="exepath" default="/opt/myapp/bin/server"/>
    ...
</server-template>
```

See Also

- IceGrid Templates
- Server-Template Descriptor Element
- Service-Template Descriptor Element
- Using Descriptor Variables and Parameters

Copyright 2016, ZeroC, Inc.

Properties Descriptor Element

The `properties` element is used in three situations:

- as a named property set if the `id` attribute is specified
- as a reference to a named property set if the `refid` attribute is specified
- as an unnamed property set if the `id` or `refid` attributes are not specified.

A property set is useful for defining a set of properties (a named property set) in application or node descriptors. Named property sets can be included in named or unnamed property sets with property set references.

A named property set element may only be a child of an `application` element or a `node` element. An unnamed property set element may only be a child of a `server`, `icebox`, `service`, `server-instance` or `service-instance` element. An unnamed property set element with the `service` attribute defined may only be a child of a `server-instance` element. A reference to a named property set can only be a child of a named or unnamed property set element.

The following attributes are supported:

Attribute	Description	Required
`id`	Defines a new named property set with the given identifier. The identifier must be unique among all named property sets defined in the same scope. If not specified, the properties element refers to an unnamed property set or a property set reference.	No
`refid`	Defines a reference to the named property set with the given identifier. If not specified, the element refers to an unnamed or named property set.	No
`service`	Specifies the name of an IceBox service that is defined in the enclosing `server-instance` descriptor. The server instance must be an IceBox server that includes a service with the given name. An unnamed property set with this attribute defined extends the properties of the service. If not specified, the unnamed property set extends the properties of the server instance.	No

Here is an example to demonstrate the use of this element:

XML

```
<application name="Simple">
   <properties id="Debug">
        <property name="Ice.Trace.Network" value="1"/>
   </properties>

    <server id="MyServer" exe="./server">
        <properties>
            <properties refid="Debug"/>
            <property name="AppProperty" value="1"/>
        </properties>
    </server>
</application>
```

See Also

- Properties and Configuration
- Application Descriptor Element
- Node Descriptor Element
- Server Descriptor Element
- IceBox Descriptor Element
- Service Descriptor Element
- Server-Instance Descriptor Element
- Service-Instance Descriptor Element

Copyright 2016, ZeroC, Inc.

Property Descriptor Element

An IceGrid node generates a configuration file for each of its servers and services. This file generally should not be edited manually because any changes are lost the next time the node generates the file. The `property` element is the correct way to define additional properties in a configuration file.

Note that IceGrid administrative utilities can retrieve the configuration properties of a server or service via the administrative facility.

This element may only appear as a child of a `server` element, a `service` element, an `icebox` element or a `properties` element.

The following attributes are supported:

Attribute	Description	Required
`name`	Specifies the property name.	Yes
`value`	Specifies the property value. If not defined, the value is an empty string.	No

Here is an example to demonstrate the use of this element:

XML

```
<server id="MyServer" ...>
    <property name="Ice.ThreadPool.Server.SizeMax" value="10"/>
    ...
</server>
```

This `property` element adds the following definition to the server's configuration file:

```
Ice.ThreadPool.Server.SizeMax=10
```

See Also

- Properties and Configuration
- icegridadmin Command Line Tool
- IceGrid and the Administrative Facility
- Server Descriptor Element
- Service Descriptor Element
- IceBox Descriptor Element
- Properties Descriptor Element

Copyright 2016, ZeroC, Inc.

Replica-Group Descriptor Element

A `replica-group` element creates a virtual object adapter in order to provide replication and load balancing for a collection of object adapters. An `adapter` element declares its membership in a group by identifying the desired replica group. The element may declare well-known objects that are available in all of the participating object adapters. A `replica-group` element may contain a `load-balancing` child element that specifies the load-balancing algorithm the registry should use when resolving locate requests. If not specified, the registry uses a random load balancing policy with the number of replicas set to 0.

This element may only appear as a child of an `application` element.

The following attributes are supported:

Attribute	Description	Required
`id`	Specifies the identifier of the replica group, which must be unique among all adapters and replica groups in the registry. This identifier can be used in indirect proxies in place of an adapter identifier.	Yes
`filter`	Specifies the identifier of a replica group filter that performs custom load balancing.	No
`proxy-options`	The proxy options to use for well-known object proxies declared with this replica group.	No

An optional nested `description` element provides free-form descriptive text.

Here is an example to demonstrate the use of this element:

XML

```
<application name="MyApp">
    <replica-group id="ReplicatedAdapter" filter="myCustomFilter" proxy-options="-e
1.0">
        <load-balancing type="adaptive" load-sample="15" n-replicas="3"/>
        <description>A description of this replica group.</description>
        <object identity="WellKnownObject" .../>
    </replica-group>
    ...
</application>
```

In this example, the proxy `WellKnownObject` is equivalent to the proxy `WellKnownObject@ReplicatedAdapter`.

See Also

- Object Adapter Replication
- Load Balancing
- Adapter Descriptor Element
- Application Descriptor Element
- Description Descriptor Element
- Well-Known Objects

Copyright 2016, ZeroC, Inc.

Server Descriptor Element

A `server` element defines a server to be deployed on a node. It typically contains at least one `adapter` element, and may supply additional information such as command-line options, environment variables, configuration properties, and a server distribution.

This element may only appear as a child of a `node` element or a `server-template` element.

The following attributes are supported:

Attribute	Description	Required
`activation`	Specifies whether the server's activation mode. Legal values are `manual`, `on-demand`, `always` and `session`. If not specified, manual activation is used by default.	No
`activation-timeout`	Defines the number of seconds a node will wait for the server to activate. If the timeout expires, a client waiting to receive the endpoints of an object adapter in this server will receive an empty set of endpoints. If not defined, the default timeout is the value of the `IceGrid.Node.WaitTime` property configured for the server's node.	No
`allocatable`	Specifies whether the server can be allocated. A server is allocated implicitly when one of its allocatable objects is allocated. The value of this attribute is ignored if the server activation mode is `session`; a server with this activation mode is always allocatable. Otherwise, if not specified and the server activation mode is not `session`, the server is not allocatable.	No
`application-distrib`	Specifies whether this server's distribution is dependent on the application's distribution. If the value is `true`, the server cannot be patched until the application has been patched. If not defined, the default value is `true`.	No
`deactivation-timeout`	Defines the number of seconds a node will wait for the server to deactivate. If the timeout expires, the node terminates the server process. If not defined, the default timeout is the value of the node's configuration property `IceGrid.Node.WaitTime`.	No
`exe`	The pathname of the server executable. On Windows, if a relative path is specified, the executable is searched in the system path. On Linux, it's searched in the user path specified by the PATH environment variable from the environment where the `icegridnode` is started.	Yes
`ice-version`	Specifies the Ice version in use by this server. If not defined, IceGrid assumes the server uses the same version that IceGrid itself uses. A server that uses an Ice version prior to 3.3 must define this attribute if its adapters use the `register-process` attribute. For example, a server using Ice 3.2.x should use `3.2` as the value of this attribute.	No
`id`	Specifies the identifier for this server. The identifier must be unique among all servers in the registry. Within the server, child elements can refer to its identifier using the reserved variable `${server}`.	Yes
`pwd`	The default working directory for the server. If not defined, the server is started in the node's current working directory.	No
`user`	Specifies the name of the user account under which the server is activated and run. If a user account mapper is configured for the node, the value of this attribute is used to find the corresponding account in the map. On Unix, the node must be running as root to be able to run servers under a different user account. On Windows, or if the node is not running as root on Unix, the only permissible value for this attribute is an empty string or the name of the user account under which the node is running. On Unix, if the node is running as root and this attribute is not specified, the server is run under the user `${session.id}` if the server activation mode is `session` or under the user `nobody` if the activation mode is `manual`, `on-demand` or `always`.	No

An optional nested `description` element provides free-form descriptive text.

Here is an example to demonstrate the use of this element:

Copyright 2016, ZeroC, Inc.

XML

```
<server id="MyServer"
        activation="on-demand"
        activation-timeout="60"
        application-distrib="false"
        deactivation-timeout="60"
        exe="/opt/app/bin/myserver"
        pwd="/">
    <option>--Ice.Trace.Network=1</option>
    <env>PATH=/opt/Ice/bin:$PATH</env>
    <property name="ServerId" value="${server}"/>
    <adapter name="Adapter1" .../>
</server>
```

See Also

- Adapter Descriptor Element
- Properties Descriptor Element
- Node Descriptor Element
- Server-Template Descriptor Element
- Description Descriptor Element
- Using Command Line Options in Descriptors
- Setting Environment Variables in Descriptors
- Application Distribution
- IceGrid Server Activation
- Using Descriptor Variables and Parameters

Copyright 2016, ZeroC, Inc.

Server-Instance Descriptor Element

A `server-instance` element deploys an instance of a `server-template` element on a node. It may supply additional information such as configuration properties.

This element may only appear as a child of a `node` element.

The following attributes are supported:

Attribute	Description	Required
`template`	Identifies the server template.	Yes

All other attributes of the element must correspond to parameters declared by the template. The `server-instance` element must provide a value for each parameter that does not have a default value supplied by the template.

Here is an example to demonstrate the use of this element:

XML

```
<icegrid>
    <application name="SampleApp">
        <server-template id="ServerTemplate">
            <parameter name="id"/>
            <server id="${id}" activation="manual" .../>
        </server-template>
        <node name="Node1">
            <server-instance template="ServerTemplate" id="TheServer">
                <properties>
                    <property name="Debug" value="1"/>
                </properties>
            </server-instance>
        </node>
    </application>
</icegrid>
```

See Also

- Server-Template Descriptor Element
- Node Descriptor Element
- IceGrid Templates
- Using Descriptor Variables and Parameters

Copyright 2016, ZeroC, Inc.

Server-Template Descriptor Element

A `server-template` element defines a template for a `server` element, simplifying the task of deploying multiple instances of the same server definition. The template should contain a parameterized `server` element that is instantiated using a `server-instance` element.

This element may only appear as a child of an `application` element.

The following attributes are supported:

Attribute	Description	Required
`id`	Specifies the identifier for the server template. This identifier must be unique among all server templates in the application.	Yes

A template may declare parameters that are used to instantiate the `server` element. You can define a default value for each parameter. Parameters without a default value are considered mandatory and values for them must be supplied by the `server-instance` element.

Here is an example to demonstrate the use of this element:

XML

```
<icegrid>
    <application name="SampleApp">
        <server-template id="ServerTemplate">
            <parameter name="id"/>
            <server id="${id}" activation="manual" .../>
        </server-template>
        <node name="Node1">
            <server-instance template="ServerTemplate" id="TheServer"/>
        </node>
    </application>
</icegrid>
```

See Also

- IceGrid Templates
- Server Descriptor Element
- Server-Instance Descriptor Element
- Application Descriptor Element
- Using Descriptor Variables and Parameters

Copyright 2016, ZeroC, Inc.

Service Descriptor Element

A `service` element defines an IceBox service. It typically contains at least one `adapter` element, and may supply additional information such as configuration properties.

This element may only appear as a child of an `icebox` element or a `service-template` element.

The following attributes are supported:

Attribute	Description	Required
`entry`	Specifies the entry point of this service.	Yes
`name`	Specifies the name of this service. Within the service, child elements can refer to its name using the reserved variable `${service}`.	Yes

An optional nested `description` element provides free-form descriptive text.

Here is an example to demonstrate the use of this element:

XML

```
<icebox id="MyIceBox" ...>
    <service name="Service1" entry="service1:Create">
        <description>A description of this service.</description>
        <property name="ServiceName" value="${service}"/>
        <adapter name="MyAdapter" id="${service}Adapter" .../>
    </service>
    <service name="Service2" entry="service2:Create"/>
</icebox>
```

See Also

- IceBox
- Adapter Descriptor Element
- Properties Descriptor Element
- IceBox Descriptor Element
- Service-Template Descriptor Element
- IceBox Integration with IceGrid

Copyright 2016, ZeroC, Inc.

Service-Instance Descriptor Element

A `service-instance` element creates an instance of a `service-template` element in an IceBox server. It may supply additional information such as configuration properties.

This element may only appear as a child of an `icebox` element.

The following attributes are supported:

Attribute	Description	Required
`template`	Identifies the service template.	Yes

All other attributes of the element must correspond to parameters declared by the template. The `service-instance` element must provide a value for each parameter that does not have a default value supplied by the template.

Here is an example to demonstrate the use of this element:

XML

```
<icebox id="IceBoxServer" ...>
    <service-instance template="ServiceTemplate" name="Service1">
        <properties>
            <property name="Debug" value="1"/>
        </properties>
    </service-instance>
</icebox>
```

See Also

- Service-Template Descriptor Element
- IceBox Descriptor Element
- IceGrid Templates
- IceBox Integration with IceGrid

Copyright 2016, ZeroC, Inc.

Service-Template Descriptor Element

A `service-template` element defines a template for a `service` element, simplifying the task of deploying multiple instances of the same service definition. The template should contain a parameterized `service` element that is instantiated using a `service-instance` element.

This element may only appear as a child of an `application` element.

The following attributes are supported:

Attribute	Description	Required
`id`	Specifies the identifier for the service template. This identifier must be unique among all service templates in the application.	Yes

A template may declare parameters that are used to instantiate the `service` element. You can define a default value for each parameter. Parameters without a default value are considered mandatory and values for them must be supplied by the `service-instance` element.

Here is an example to demonstrate the use of this element:

XML

```
<icegrid>
    <application name="IceBoxApp">
        <service-template id="ServiceTemplate">
            <parameter name="name"/>
            <service name="${name}" entry="DemoService:create">
                <adapter name="${service}" .../>
            </service>
        </service-template>
        <node name="Node1">
            <icebox id="IceBoxServer" ...>
                <service-instance template="ServiceTemplate" name="Service1"/>
            </icebox>
        </node>
    </application>
</icegrid>
```

See Also

- IceGrid Templates
- Service Descriptor Element
- Service-Instance Descriptor Element
- IceBox Integration with IceGrid
- Using Descriptor Variables and Parameters

Copyright 2016, ZeroC, Inc.

Variable Descriptor Element

A `variable` element defines a variable.

This element may only appear as a child of an `application` element or `node` element.

The following attributes are supported:

Attribute	Description	Required
`name`	Specifies the variable name. The value of this variable is substituted whenever its name is used in variable syntax, as in `${name}`.	Yes
`value`	Specifies the variable value. If not defined, the default value is an empty string.	No

Here is an example to demonstrate the use of this element:

XML

```
<icegrid>
    <application name="SampleApp">
        <variable name="Var1" value="foo"/>
        <variable name="Var2" value="${Var1}bar"/>
        ...
    </application>
</icegrid>
```

See Also

- Using Descriptor Variables and Parameters
- Application Descriptor Element
- Node Descriptor Element

Copyright 2016, ZeroC, Inc.

Using Command Line Options in Descriptors

Server descriptors and IceBox descriptors may specify command-line options that the node will pass to the program at startup. As the node prepares to execute the server, it assembles the command by appending options to the server executable's pathname.

In XML, you define a command-line option using the `option` element:

XML

```
<server id="Server1" ...>
    <option>--Ice.Trace.Protocol</option>
    ...
</server>
```

The node preserves the order of options, which is especially important for Java servers. For example, JVM options must appear before the class name, as shown below:

XML

```
<server id="JavaServer" exe="java" ...>
    <option>-Xnoclassgc</option>
    <option>ServerClassName</option>
    <option>--Ice.Trace.Protocol</option>
    ...
</server>
```

The node translates these options into the following command:

```
java -Xnoclassgc ServerClassName --Ice.Trace.Protocol
```

See Also

- Server Descriptor Element
- IceBox Descriptor Element

Copyright 2016, ZeroC, Inc.

Setting Environment Variables in Descriptors

Server descriptors and IceBox descriptors may specify environment variables that the node will define when starting a server. An environment variable definition uses the familiar *name=value* syntax, and you can also refer to other environment variables within the value. The exact syntax for variable references depends on the platform on which the server's descriptor is deployed.

On a Unix platform, the Bourne shell syntax is required:

```
LD_LIBRARY_PATH=/opt/Ice/lib:$LD_LIBRARY_PATH
```

On a Windows platform, the syntax uses the conventional style:

```
PATH=C:\Ice\lib;%PATH%
```

In XML, the `env` element supplies a definition for an environment variable:

XML

```
<node name="UnixBox">
    <server id="UnixServer" exe="/opt/app/bin/server" ...>
        <env>LD_LIBRARY_PATH=/opt/Ice/lib:$LD_LIBRARY_PATH</env>
        ...
    </server>
</node>
<node name="WindowsBox">
    <server id="WindowsServer" exe="C:/app/bin/server.exe" ...>
        <env>PATH=C:\Ice\lib;%PATH%</env>
        ...
    </server>
</node>
```

If a value refers to an environment variable that is not defined, the reference is substituted with an empty string.

Environment variable definitions may also refer to descriptor variables and template parameters:

XML

```
<node name="UnixBox">
    <server id="UnixServer" exe="/opt/app/bin/server" ...>
        <env>PATH=${server.distrib}/bin:$PATH</env>
        ...
    </server>
</node>
```

On Unix, an environment variable `VAR` can be referenced as `$VAR` or `${VAR}`. You must be careful when using the latter syntax because IceGrid assumes `${VAR}` refers to a descriptor variable or parameter and will report an error if no match is found. If you prefer to use this style to refer to environment variables, you must escape these occurrences as shown in the example below:

Copyright 2016, ZeroC, Inc.

```
<node name="UnixBox">
    <server id="UnixServer" exe="/opt/app/bin/server" ...>
        <env>PATH=${server.distrib}/bin:$${PATH}</env>
        ...
    </server>
</node>
```

IceGrid does not attempt to perform substitution on `$${PATH}`, but rather removes the leading `$` character and then performs environment variable substitution on `$${PATH}`.

See Also

- Server Descriptor Element
- IceBox Descriptor Element
- Using Descriptor Variables and Parameters

Copyright 2016, ZeroC, Inc.

Using Descriptor Variables and Parameters

Variable descriptors allow you to define commonly-used information once and refer to them symbolically throughout your application descriptors.

On this page:

- Descriptor Substitution Syntax
  - Limitations
  - Escaping a Variable
- Special Descriptor Variables
- Descriptor Variable Scoping Rules
  - Resolving a Reference
  - Template Parameters
  - Modifying a Variable

Descriptor Substitution Syntax

Substitution for a variable or parameter *VP* is attempted whenever the symbol ${*VP*} is encountered, subject to the limitations and rules described below. Substitution is case-sensitive, and a fatal error occurs if *VP* is not defined.

Limitations

Substitution is only performed in string values, and excludes the following cases:

- Identifier of a template descriptor definition

```
<server-template id="${invalid}" ...>
```

- Name of a variable definition

```
<variable name="${invalid}" ...>
```

- Name of a template parameter definition

```
<parameter name="${invalid}" ...>
```

- Name of a template parameter assignment

```
<server-instance template="T" ${invalid}="val" ...>
```

- Name of a node definition

```
<node name="${invalid}" ...>
```

- Name of an application definition

```
<application name="${invalid}" ...>
```

Substitution is not supported for values of other types. The example below demonstrates an invalid use of substitution:

Copyright 2016, ZeroC, Inc.

```
<variable name="register" value="true"/>
<node name="Node">
    <server id="Server1" ...>
        <adapter name="Adapter1" register-process=${register} .../>
```

In this case, a variable cannot supply the value of `register-process` because that attribute expects a boolean value, not a string.

Most values are strings, however, so this limitation is rarely a problem.

Escaping a Variable

You can prevent substitution by escaping a variable reference with an additional leading $ character. For example, in order to assign the literal string `${abc}` to a variable, you must escape it as shown below:

```
<variable name="x" value="$${abc}"/>
```

The extra $ symbol is only meaningful when immediately preceding a variable reference, therefore text such as `US$$55` is not modified. Each occurrence of the characters $$ preceding a variable reference is replaced with a single $ character, and that character does not initiate a variable reference. Consider these examples:

```
<variable name="a" value="hi"/>
<variable name="b" value="$${a}"/>
<variable name="c" value="$$${a}"/>
<variable name="d" value="$$$${a}"/>
```

After substitution, b has the value `${a}`, c has the value `$hi`, and d has the value `$${a}`.

Special Descriptor Variables

IceGrid defines a set of read-only variables to hold information that may be of use to descriptors. The names of these variables are reserved and cannot be used as variable or parameter names. The table describes the purpose of each variable and defines the context in which it is valid.

Reserved Name	Description
`application`	The name of the enclosing application.
`application.distrib`	The pathname of the enclosing application's distribution directory, and an alias for `${node.datadir}/distrib/${application}`.
`node`	The name of the enclosing node.
`node.os`	The name of the enclosing node's operating system. On Unix, this is value is provided by `uname`. On Windows, the value is `Windows`.
`node.hostname`	The host name of the enclosing node.
`node.release`	The operating system release of the enclosing node. On Unix, this value is provided by `uname`. On Windows, the value is obtained from the `OSVERSIONINFO` data structure.
`node.version`	The operating system version of the enclosing node. On Unix, this value is provided by `uname`. On Windows, the value represents the current service pack level.

Copyright 2016, ZeroC, Inc.

`node.machine`	The machine hardware name of the enclosing node. On Unix, this value is provided by `uname`. On Windows, the value can be x86, x64, or IA64, depending on the machine architecture.
`node.datadir`	The absolute pathname of the enclosing node's data directory.
`server`	The ID of the enclosing server.
`server.distrib`	The pathname of the enclosing server's distribution directory, and an alias for `${node.datadir}/servers/${server}/distrib`.
`service`	The name of the enclosing service.
`session.id`	The client session identifier. For sessions created with a user name and password, the value is the user ID; for sessions created from a secure connection, the value is the distinguished name associated with the connection.

The availability of a variable is easily determined in some cases, but may not be readily apparent in others. For example, the following example represents a valid use of the `${node}` variable:

XML

```
<icegrid>
    <application name="App">
        <server-template id="T" ...>
            <parameter name="id"/>
            <server id="${id}" ...>
                <property name="NodeName" value="${node}"/>
                ...
            </server>
        </server-template>
        <node name="TheNode">
            <server-instance template="T" id="TheServer"/>
        </node>
    </application>
</icegrid>
```

Although the server template descriptor is defined as a child of an application descriptor, its variables are not evaluated until it is instantiated. Since a template *instance* is always enclosed within a node, it is able to use the `${node}` variable.

Descriptor Variable Scoping Rules

Descriptors may only define variables at the application and node levels. Each node introduces a new scope, such that defining a variable at the node level overrides (but does not modify) the value of an application variable with the same name. Similarly, a template parameter overrides the value of a variable with the same name in an enclosing scope. A descriptor may refer to a variable defined in any enclosing scope, but its value is determined by the nearest scope. The following figure illustrates these concepts:

Copyright 2016, ZeroC, Inc.

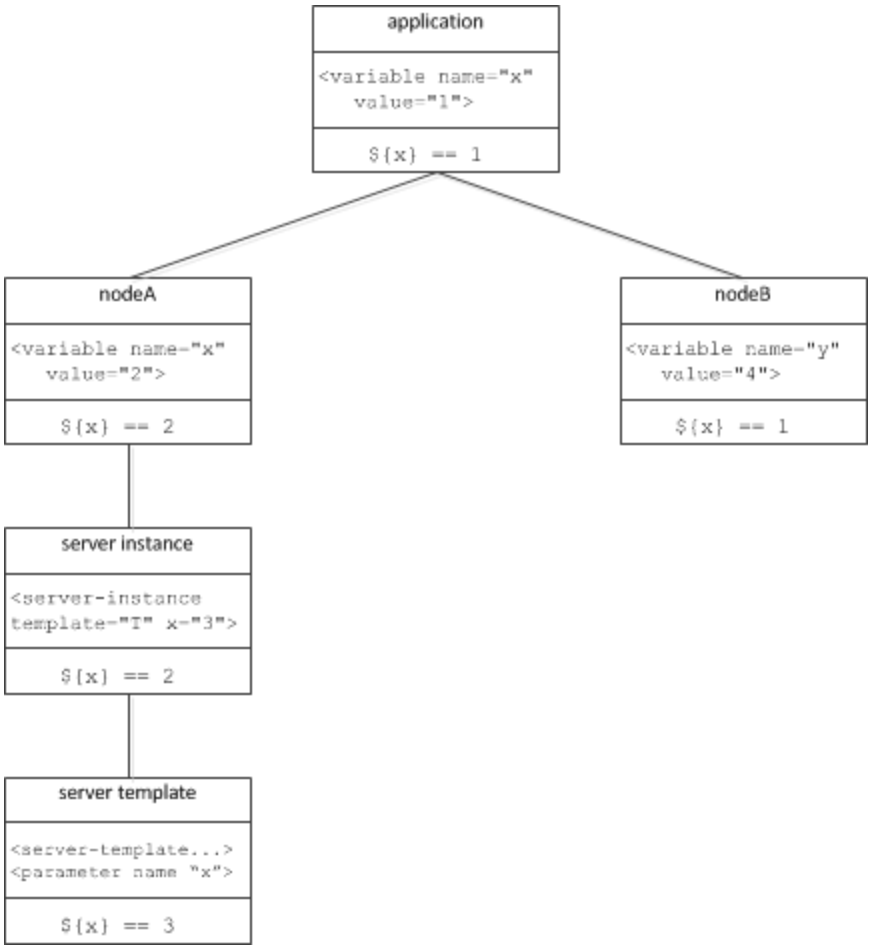

Variable scoping semantics.

In this diagram, the variable x is defined at the application level with the value 1. In nodeA, x is overridden with the value 2, whereas x remains unchanged in nodeB. Within the context of nodeA, x continues to have the value 2 in a server instance definition. However, when x is used as the name of a template parameter, the node's definition of x is overridden and x has the value 3 in the template's scope.

Resolving a Reference

To resolve a variable reference `${var}`, IceGrid searches for a definition of `var` using the following order of precedence:

1. Pre-defined variables
2. Template parameters, if applicable
3. Node variables, if applicable
4. Application variables

After the initial substitution, any remaining references are resolved recursively using the following order of precedence:

1. Pre-defined variables
2. Node variables, if applicable
3. Application variables

Template Parameters

Template parameters are not visible in nested template instances. This situation can only occur when an IceBox server template instantiates a service template, as shown in the following example:

Copyright 2016, ZeroC, Inc.

XML

```
<icegrid>
    <application name="IceBoxApp">
        <service-template id="ServiceTemplate">
            <parameter name="name"/>
            <service name="${name}" entry="DemoService:create">
                ...
                <property name="${name}.Identity"
value="${id}-${name}"/> <!-- WRONG! -->
            </service>
        </service-template>
        <server-template id="ServerTemplate">
            <parameter name="id"/>
            <icebox id="${id}" endpoints="default" ...>
                <service-instance template="ServiceTemplate" name="Service1"/>
            </icebox>
        </server-template>
        <node name="Node1">
            <server-instance template="ServerTemplate" id="IceBoxServer"/>
        </node>
    </application>
</icegrid>
```

The service template incorrectly refers to `id`, which is a parameter of the server template.

Template parameters can be referenced only in the body of a template; they cannot be used to define other parameters. For example, the following is illegal:

XML

```
<server-template id="ServerTemplate">
    <parameter name="par1"/>
    <parameter name="par2" default="${par1}"/>
    ...
</server-template>
```

Modifying a Variable

A variable definition can be overridden in an inner scope, but the inner definition does not modify the outer variable. If a variable is defined multiple times in the same scope (which is only relevant in XML definitions), the most recent definition is used for all references to that variable. Consider the following example:

Copyright 2016, ZeroC, Inc.

XML

```
<application name="MyApp">
    <variable name="x" value="1"/>
    <variable name="y" value="${x}"/>
    <variable name="x" value="2"/>
    ...
</application>
```

When descriptors such as these are created, IceGrid validates their variable references but does not perform substitution until the descriptor is acted upon (such as when a node is generating a configuration file for a server). As a result, the value of y in the above example is 2 beca use that is the most recent definition of x.

See Also

- Variable Descriptor Element
- IceGrid Templates
- Application Distribution
- Variables in IceGrid Descriptors

Copyright 2016, ZeroC, Inc.

IceGrid Property Set Semantics

Ice servers and clients are configured with properties. For servers deployed with IceGrid, these properties are automatically generated into a configuration file from the information contained in the application descriptor. The settings in that configuration file are passed to server via the `--Ice.Config` command-line option.

Property descriptors allow you to define property sets to efficiently manage and specify properties. Here are some of the benefits of using property sets:

- You can define sets of properties at the application or node element level and reference these properties in other property sets.
- You can specify properties for a specific server or service instance.

There are two kinds of property sets:

- Named property sets
  Named property sets are defined at the application or node level. They are useful only as the target of references from other property sets. Specifically, a named property set has no effect unless you reference it from a server descriptor.
- Unnamed property sets
  Unnamed property sets can be defined in server, service, icebox, server-instance or service-instance elements and define the properties for a server or service. Unnamed property sets can reference named property sets.

Named and unnamed property sets are defined with the same properties descriptor. The context and the attributes of a `properties` element distinguish named property sets from unnamed property sets. Here is an example that defines a named and an unnamed property set:

XML

```
<application name="App">
    <properties id="Debug">
        <property name="UseDebug" value="1"/>
    </properties>

    <node name="TheNode">
        <server id="TheServer" exe="./server">
            <properties>
                <property name="Identity" value="hello"/>
            </properties>
        </server>
    </node>
</application>
```

In this example, we define the named property set `Debug` and the unnamed property set of the server `TheServer`. The server configuration will contain only the `Identity` property because the server property set does not reference the `Debug` named property set.

The `properties` element is used to reference a named property set: if a `properties` element appears inside another `properties` element, it is a reference to another property set and it must specify the `refid` attribute. With the previous example, to reference the `Debug` property set, we would write the following:

Copyright 2016, ZeroC, Inc.

XML

```
<application name="App">
    <properties id="Debug">
        <property name="UseDebug" value="1"/>
    </properties>

    <node name="TheNode">
        <server id="TheServer" exe="./server">
            <properties>
                <properties refid="Debug"/>
                <property name="Identity" value="hello"/>
        </properties>
        </server>
    </node>
</application>
```

Property sets, whether named or unnamed, are evaluated as follows:

1. Within a `properties` element, IceGrid locates all references to named property sets and evaluates all property settings in the referenced property sets.
2. Explicit property definitions following any named references are then evaluated and added to the property set formed in the preceding step. This means that explicit property settings override corresponding settings in any referenced property sets.

It is illegal to define a reference to a property set after setting a property value, so references to property sets must precede property definitions. For example, the following is illegal:

XML

```
<properties>
    <property name="Prop1" value="Value1"/>
    <properties refid="Ref1"/>
</properties>
```

Just as the order of the property definitions is important, the order of property set references is also important. For example, the following two property sets are not equivalent:

XML

```
<properties>
    <properties refid="Ref1"/>
    <properties refid="Ref2"/>
</properties>

<properties>
    <properties refid="Ref2"/>
    <properties refid="Ref1"/>
</properties>
```


Copyright 2016, ZeroC, Inc.

Named property sets are evaluated at the point of definition. If you reference other property sets or use variables in a named property set definition, you must make sure that the referenced property sets or variables are defined in the same scope. For example, the following is correct:

XML

```
<application name="App">

    <variable name="level" value="1"/>

    <properties id="DebugApp">
        <property name="DebugLevel value="${level}">
    </properties>

</application>
```

However, the following example is wrong because the `${level}` variable is not defined at the `application` scope:

XML

```
<application name="App">

    <properties id="DebugApp">
        <property name="DebugLevel value="${level}">
    </properties>

    <node name="TheNode">
        <variable name="level" value="1"/>
    </node>

</application>
```

If both the `application` and the `node` define the `${level}` variable, the value of the `${level}` variable in the `DebugApp` property set will be the value of the variable defined in the application descriptor.

So far, we have seen the definition of an unnamed property set only in a server descriptor. However, it is also possible to define an unnamed property set for server or service instances. This is a good way to specify or override properties specific to a server or service instance. For example:

Copyright 2016, ZeroC, Inc.

XML

```
<application name="TheApp">
  <server-template id="Template">

    <parameter name="instance-name"/>

    <server id="${instance-name}" exe="./server">
      <properties>
        <property name="Timeout" value="30"/>
      </properties>
    </server>
  </server-template>

  <node name="TheNode">
    <server-instance template="Template" instance-name="MyInst">
      <properties>
        <property name="Debug" value="1"/>
        <property name="Timeout" value="-1"/>
      </properties>
    </server-instance>
  </node>
</application>
```

Here, the server instance overrides the `Timeout` property and defines an additional `Debug` property.

The server or service instance properties are evaluated as follows:

1. The unnamed property set from the template server or service descriptor is evaluated.
2. The unnamed property set from the server or service instance descriptor is evaluated and the resulting properties are added to the property set formed in the preceding step. This means that property settings in a server or service instance descriptor override corresponding settings in a template server or service descriptor.

The server or service instance unnamed property set and its parameters provide two different ways to customize the properties of a server or service template instance. It might not always be obvious which method to use: is it better to use a parameter to parameterize a given property or is it better to just specify it in the server or service instance property set?

For example, in the previous descriptor, we could have used a parameter with a default value for the `Timeout` property:

Copyright 2016, ZeroC, Inc.

XML

```
<application name="TheApp">
  <server-template id="Template">

    <parameter name="instance-name"/>
    <parameter name="timeout" default="30"/>

    <server id="${instance-name}" exe="./server">
      <properties>
        <property name="Timeout" value="${timeout}"/>
      </properties>
    </server>
  </server-template>

  <node name="TheNode">
    <server-instance template="Template" instance-name="MyInst" timeout="-1">
      <properties>
        <property name="Debug" value="1"/>
      </properties>
    </server-instance>
  </node>
</application>
```

Here are some guidelines to help you decide whether to use a parameter or a property:

- Use a parameter for a property that should always be set.
- Use a parameter if you want to make the property obvious to the reader and user of the template.
- Do not use a parameter for optional properties if you want to rely on a default value for the server.
- Do not use parameters for properties that are rarely used.

See Also

- Properties and Configuration
- Using IceGrid Deployment
- Properties Descriptor Element
- Application Descriptor Element
- Node Descriptor Element
- Server-Instance Descriptor Element
- Service-Instance Descriptor Element

Copyright 2016, ZeroC, Inc.

IceGrid XML Features

IceGrid provides some convenient features to simplify the task of defining descriptors in XML.

On this page:

- Adding Flexibility with Targets
- Including Descriptor Files

Adding Flexibility with Targets

An IceGrid XML file may contain optional definitions that are deployed only when specifically requested. These definitions are called targets and must be defined within a `target` element. The elements that may legally appear within a `target` element are determined by its enclosing element. For example, a `node` element is legal inside a `target` element of an `application` element, but not inside a `target` el ement of a `server` element. Each `target` element must define a value for the `name` attribute, but names are not required to be unique. Rather, targets should be considered as optional components or features of an application that are deployed in certain circumstances.

The example below defines targets named `debug` that, if requested during deployment, configure their servers with an additional property:

XML

```
<icegrid>
    <application name="MyApp">
        <node name="Node">
            <server id="Server1" ...>
                <target name="debug">
                    <property name="Ice.Trace.Network" value="2"/>
                </target>
                ...
            </server>
            <server id="Server2" ...>
                <target name="debug">
                    <property name="Ice.Trace.Network" value="2"/>
                </target>
                ...
            </server>
        </node>
    </application>
</icegrid>
```

Target names specified in an `icegridadmin` command can be unqualified names like `debug`, in which case every target with that name is deployed, regardless of the target's nesting level. If you want to deploy targets more selectively, you can specify a fully-qualified name instead. A fully-qualified target name consists of its unqualified name prefaced by the names or identifiers of each enclosing element. For instance, a fully-qualified target name from the example above is `MyApp.Node.Server1.debug`.

Including Descriptor Files

You can include the contents of another XML file into the current file using the `include` element, which is replaced with the contents of the included file. The elements in the included file must be enclosed in an `icegrid` element, as shown in the following example:

Copyright 2016, ZeroC, Inc.

XML

```
<!-- File: A.xml -->
<icegrid>
    <server-template id="ServerTemplate">
        <parameter name="id"/>
        ...
    </server-template>
</icegrid>

<!-- File: B.xml -->
<icegrid>
    <application name="MyApp">
        <include file="A.xml"/>
        <node name="Node">
            <server-instance template="ServerTemplate" .../>
        </node>
    </application>
</icegrid>
```

In `B.xml`, the `include` element identifies the name of the file to include using the `file` attribute. The top-level `icegrid` element is discarded from `A.xml` and its contents are inserted at the position of the `include` element in `B.xml`.

Note that the file name of an included file is relative to the application descriptor, not relative to the working directory.

You can include specific targets from a file by specifying their names in the optional `targets` attribute. If multiple targets are included, their names must be separated by whitespace. The example below illustrates the use of a target:

Copyright 2016, ZeroC, Inc.

XML

```
<!-- File: A.xml -->
<icegrid>
    <server-template id="ServerTemplate">
        <parameter name="id"/>
        ...
    </server-template>
    <target name="targetA">
        <server-template id="AnotherTemplate">
            ...
        </server-template>
    </target>
</icegrid>

<!-- File: B.xml -->
<icegrid>
    <application name="MyApp">
        <include file="A.xml" targets="targetA"/>
        <node name="Node">
            <server-instance template="ServerTemplate" .../>
            <server-instance template="AnotherTemplate" .../>
        </node>
    </application>
</icegrid>
```

See Also

- Using IceGrid Deployment
- icegridadmin Command Line Tool

Copyright 2016, ZeroC, Inc.

IceGrid Server Reference

Topics

- icegridregistry
- icegridnode
- Well-Known Registry Objects
- IceGrid Persistent Data
- Promoting a Registry Slave

Copyright 2016, ZeroC, Inc.

icegridregistry

The IceGrid registry is a centralized repository of information, including deployed applications and well-known objects. A registry can optionally be collocated with an IceGrid node, which conserves resources and can be convenient during development and testing. The registry server is implemented by the `icegridregistry` executable.

On this page:

- Command Line Options for icegridregistry
- Configuring Registry Endpoints
- Registry Security Considerations
- Configuring a Data Directory for the Registry
- Registry Configuration Example

Command Line Options for `icegridregistry`

The registry supports the following command-line options:

```
$ icegridregistry -h
Usage: icegridregistry [options]
Options:
-h, --help           Show this message.
-v, --version        Display the Ice version.
--nowarn             Don't print any security warnings.
--readonly           Start the master registry in read-only mode.
--initdb-from-replica=<replica>
                     Initialize the database from the given replica.
```

The `--readonly` option prevents any updates to the registry's database; it also prevents slaves from synchronizing their databases with this master. This option is useful when you need to verify that the master registry's database is correct after promoting a slave to become the new master. The `--initdb-from-replica` option, added in Ice 3.5.1, allows you to initialize the database from another registry replica. This option is useful when you need to start a new master with the contents of a slave database.

Additional command line options are supported, including those that allow the registry to run as a Windows service or Unix daemon, and Ice includes a utility to help you install an IceGrid registry as a Windows service.

Configuring Registry Endpoints

The IceGrid registry creates up to five sets of endpoints, configured with the following properties:

- `IceGrid.Registry.Client.Endpoints`
  Client-side endpoints supporting the following interfaces:
  - `Ice::Locator`
  - `IceGrid::Query`
  - `IceGrid::Registry`
  - `IceGrid::Session`
  - `IceGrid::AdminSession`
  - `IceGrid::Admin`

  There are security implications in allowing access to administrative sessions, as explained in the next section.
- `IceGrid.Registry.Server.Endpoints`
  Server-side endpoints for object adapter registration.
- `IceGrid.Registry.SessionManager.Endpoints`
  Optional endpoints for delegating session authentication to a Glacier2 router.
- `IceGrid.Registry.AdminSessionManager.Endpoints`
  Optional endpoints for delegating administrative session authentication to a Glacier2 router.
- `IceGrid.Registry.Internal.Endpoints`

Copyright 2016, ZeroC, Inc.

Internal endpoints used by IceGrid nodes and registry replicas. This property must be defined even if no nodes or replicas are being used.

Registry Security Considerations

A client that successfully establishes an administrative session with the registry has the ability to compromise the security of the registry host. As a result, it is imperative that you configure the registry carefully if you intend to allow the use of administrative sessions.

Administrative sessions are disabled unless you explicitly configure the registry to use an authentication mechanism. To allow authentication with a user name and password, you can specify a password file using the property `IceGrid.Registry.AdminCryptPasswords` or use your own permissions verifier object by supplying its proxy in the property `IceGrid.Registry.AdminPermissionsVerifier`. Your verifier object must implement the `Glacier2::PermissionsVerifier` interface.

To authenticate administrative clients using their SSL connections, define `IceGrid.Registry.AdminSSLPermissionsVerifier` with the proxy of a verifier object that implements the `Glacier2::SSLPermissionsVerifier` interface.

Configuring a Data Directory for the Registry

You must provide an empty directory in which the registry can initialize its databases. The path name of this directory is supplied by the configuration property `IceGrid.Registry.Data`.

The files in this directory must not be edited manually, but rather indirectly using one of the administrative tools. To clear a registry's databases, first ensure the server is not currently running, then remove all of the files in its data directory and restart the server.

Registry Configuration Example

The registry requires values for the three mandatory endpoint properties, as well as the data directory property, as shown in the following example:

```
IceGrid.Registry.Client.Endpoints=tcp -p 4061
IceGrid.Registry.Server.Endpoints=tcp
IceGrid.Registry.Internal.Endpoints=tcp
IceGrid.Registry.Data=/opt/ripper/registry
```

In addition, we also recommend defining `IceGrid.InstanceName`, whose value affects the identities of the registry's well-known objects.

The remaining configuration properties are discussed in IceGrid.*.

See Also

- Using IceGrid Deployment
- Well-Known Objects
- Promoting a Registry Slave
- Windows Services
- Resource Allocation using IceGrid Sessions
- IceGrid Administrative Sessions
- Glacier2 Integration with IceGrid
- icegridadmin Command Line Tool
- Securing a Glacier2 Router
- IceGrid Persistent Data
- Well-Known Registry Objects
- IceGrid.*

Copyright 2016, ZeroC, Inc.

icegridnode

An IceGrid node is a process that activates, monitors, and deactivates registered server processes. You can run any number of nodes in a domain, but typically there is one node per host. A node must be running on each host on which servers are activated automatically, and nodes cannot run without an IceGrid registry.

The IceGrid node server is implemented by the `icegridnode` executable. If you wish to run a registry and node in one process, `icegridnode` is the executable you must use.

On this page:

- Command Line Options for icegridnode
- Configuring Node Endpoints
- Node Security Considerations
- Configuring a Data Directory for the Node
- Node Configuration Example

Command Line Options for `icegridnode`

The node supports the following command-line options:

```
Usage: icegridnode [options]
Options:
-h, --help           Show this message.
-v, --version        Display the Ice version.
--nowarn             Don't print any security warnings.
--readonly           Start the collocated master registry in
                     read-only mode.
--initdb-from-replica=<replica>
                     Initialize the collocated registry database from the
                     given replica.

--deploy DESCRIPTOR [TARGET1 [TARGET2 ...]]
                     Add or update descriptor in file DESCRIPTOR,
                     with optional targets.
```

If you are running the node with a collocated registry, the `--readonly` option prevents any updates to the registry's database; it also prevents slaves from synchronizing their databases with this master. This option is useful when you need to verify that the master registry's database is correct after promoting a slave to become the new master. The `--initdb-from-replica` option, added in Ice 3.5.1, allows you to initialize the database from another registry replica. This option is useful when you need to start a new master with the contents of a slave database.

The `--deploy` option allows an application to be deployed automatically as the node process starts, which can be especially useful during testing. The command expects the name of the XML deployment file, and optionally allows the names of the individual targets within the file to be specified.

Additional command line options are supported, including those that allow the node to run as a Windows service or Unix daemon, and Ice includes a utility to help you install an IceGrid node as a Windows service.

Configuring Node Endpoints

The IceGrid node's endpoints are defined by the `IceGrid.Node.Endpoints` property and must be accessible to the registry. It is not necessary to use a fixed port because each node contacts the registry at startup to provide its current endpoint information.

Node Security Considerations

It is important that you give careful consideration to the permissions of the account under which the node runs. If the servers that the node will activate have no special access requirements, and all of the servers can use the same account, it is recommended that you do not run the node under an account with system privileges, such as the root account on Unix or the Administrator account on Windows.

Copyright 2016, ZeroC, Inc.

Configuring a Data Directory for the Node

The node requires an empty directory that it can use to store server files. The pathname of this directory is supplied by the configuration property `IceGrid.Node.Data`. To clear a node's state, first ensure the server is not currently running, then remove all of the files in its data directory and restart the server.

The node's data directory may also contain files and subdirectories used by your application's servers, such as configuration files and Freeze database environments. Before destroying the contents of the node's data directory, make sure that all servers are stopped and any important files are backed up.

When running a collocated node and registry server, we recommend using separate directories for the registry and node data directories.

Node Configuration Example

A minimal node configuration is shown in the following example:

```
IceGrid.Node.Endpoints=tcp
IceGrid.Node.Name=Node1
IceGrid.Node.Data=/opt/ripper/node

Ice.Default.Locator=IceGrid/Locator:tcp -p 4061
```

The value of the `IceGrid.Node.Name` property must match that of a deployed node known by the registry.

The `Ice.Default.Locator` property is used by the node to contact the registry. The value is a proxy that contains the registry's client endpoints. (You can avoid the need to define `Ice.Default.Locator` by using IceLocatorDiscovery.)

If you wish to run a collocated registry and node server, enable the property `IceGrid.Node.CollocateRegistry` and include the registry's configuration properties.

The remaining configuration properties are discussed in IceGrid.*.

See Also

- IceGrid Server Activation
- Promoting a Registry Slave
- IceGrid Persistent Data
- Getting Started with IceGrid
- Using IceGrid Deployment
- Windows Services
- IceGrid.*

Copyright 2016, ZeroC, Inc.

Well-Known Registry Objects

On this page:

- Default Identities of Registry Objects
- Using the LocatorFinder Interface

Default Identities of Registry Objects

The IceGrid registry hosts several well-known objects. The following table shows the default identities of these objects and their corresponding Slice interfaces:

Default Identity	Interface
`IceGrid/AdminSessionManager`	`Glacier2::SessionManager`
`IceGrid/AdminSessionManager-`*`replica`*	`Glacier2::SessionManager`
`IceGrid/AdminSSLSessionManager`	`Glacier2::SSLSessionManager`
`IceGrid/AdminSSLSessionManager-`*`replica`*	`Glacier2::SSLSessionManager`
`IceGrid/Locator`	`Ice::Locator`
`IceGrid/Query`	`IceGrid::Query`
`IceGrid/Registry`	`IceGrid::Registry`
`IceGrid/Registry-`*`replica`*	`IceGrid::Registry`
`IceGrid/RegistryUserAccountMapper`	`IceGrid::UserAccountMapper`
`IceGrid/RegistryUserAccountMapper-`*`replica`*	`IceGrid::UserAccountMapper`
`IceGrid/SessionManager`	`Glacier2::SessionManager`
`IceGrid/SSLSessionManager`	`Glacier2::SSLSessionManager`
`Ice/LocatorFinder`	`Ice::LocatorFinder`

You can assign unique identities to these objects by configuring the `IceGrid.InstanceName` property, as shown in the following example:

```
IceGrid.InstanceName=MP3Grid
```

This property changes the identities of the well-known objects to use `MP3Grid` instead of `IceGrid` as the identity category. For example, the identity of the locator becomes `MP3Grid/Locator`.

Changes to `IceGrid.InstanceName` do not affect the identity of the `Ice/LocatorFinder` object. See the next section for more information.

The client's configuration must also be changed to reflect the new identity:

```
Ice.Default.Locator=MP3Grid/Locator:tcp -h registryhost -p 4061
```

Furthermore, any uses of these identities in application code must be updated as well.

Using the `LocatorFinder` Interface

Copyright 2016, ZeroC, Inc.

As of Ice 3.6, Ice requires all locator implementations to support the `Ice::LocatorFinder` interface:

Slice

```
module Ice {
interface LocatorFinder {
    Locator* getLocator();
};
};
```

An object supporting this interface must be available with the identity `Ice/LocatorFinder`. By knowing the host and port of a locator's client endpoints, a client can discover the locator's proxy at run time with a call to `getLocator`:

C++

```
Ice::ObjectPrx prx = communicator->stringToProxy("Ice/LocatorFinder:tcp -p 4061 -h
gridhost");
Ice::LocatorFinderPrx finder = Ice::LocatorFinderPrx::checkedCast(prx);
Ice::LocatorPrx locator = finder->getLocator();
communicator->setDefaultLocator(locator);
```

See Also

- Well-Known Objects

Copyright 2016, ZeroC, Inc.

IceGrid Persistent Data

The IceGrid registry and node both store information in the data directories specified by the `IceGrid.Registry.Data` and `IceGrid.Node.Data` properties, respectively. This section describes what the registry and node are storing and discusses backup and recovery techniques.

On this page:

- Registry Persistent Data
- Node Persistent Data

Registry Persistent Data

The registry's data directory is a Freeze database environment that stores the following information:

- Applications deployed using the `addApplication` operation on the `IceGrid::Admin` interface (which includes the IceGrid GUI and command-line administrative clients). Applications specify servers, well-known objects, object adapters, replica groups, and allocatable objects. Applications can be removed with the `removeApplication` operation.
- Well-known objects registered using the `addObject` and `addObjectWithType` operations on the `IceGrid::Admin` interface. Well-known objects added by these operations can be removed using the `removeObject` operation.
- Adapter endpoints registered dynamically by servers using the `Ice::LocatorRegistry` interface. The property `IceGrid.Registry.DynamicRegistration` must be set to a value larger than zero to allow the dynamic registration of object adapters. These adapters can be removed using the `removeAdapter` operation.
- Some internal proxies used by the registry to contact nodes and other registry replicas during startup. The proxies enable the registry to notify these entities about the registry's availability.

Client session and administrative sessions established with the IceGrid registry are not persistent. If the registry is restarted, these sessions must be recreated. For client sessions in particular, this implies that objects allocated using the allocation mechanism will no longer be allocated once the IceGrid registry restarts.

If the registry's database is corrupted or lost, you must recover the deployed applications, the well-known objects, and the adapter endpoints. You do not need to worry about the internal proxies stored by the registry, as the nodes and registry replicas will eventually contact the registry again.

Depending on your deployed applications and your use of the registry, you should consider backing up the registry's database, especially if you cannot easily recover the persistent information.

For example, if you rely on dynamically-registered adapters, or on well-known objects registered programmatically via the `IceGrid::Admin` interface, you should back up the registry database because recovering this information may be difficult. On the other hand, if you only deploy a few applications from XML files, you can easily recover the applications by redeploying their XML files, and therefore backing up the database may be unnecessary.

Be aware that restarting the registry with an empty database may cause the server information stored by the nodes to be deleted. This can be an issue if the deployed servers have databases stored in the node data directory. The next section provides more information on this subject.

Node Persistent Data

The IceGrid node stores information for servers deployed by IceGrid applications. This information is stored in the `servers` subdirectory of the node's data directory. There is one subdirectory per server; the name of the subdirectory is the server's ID. Each server directory contains configuration files, and may also contain database environments and the distribution data of the server. The node's data directory also contains a `distrib` directory to store per-application distribution data. This directory contains a subdirectory for each application that specifies a distribution and has a server deployed on the node.

If a server directory is deleted, the node recreates it at startup. The node will also recreate the server configuration files and the database environment directories. However, the node cannot restore the prior contents of a server's database environment. It is your responsibility to back up these database environments and restore them when necessary. If the server or application distribution data is deleted from the node's data directory, you can easily recover the deleted information by patching these distributions again using the IceGrid administrative tools.

If you store your server database environments outside the node's data directory (such as in a directory that is regularly backed up), or if you do not have any database environments inside the node's data directory, you do not need to back up the contents of the node's data directory.

See Also

- Freeze

Copyright 2016, ZeroC, Inc.

- Using IceGrid Deployment
- icegridadmin Command Line Tool
- Well-Known Objects
- Object Adapter Endpoints
- Resource Allocation using IceGrid Sessions
- IceGrid Administrative Sessions
- DbEnv Descriptor Element
- Application Distribution
- IceGrid.*

Copyright 2016, ZeroC, Inc.

Promoting a Registry Slave

In a replicated IceGrid deployment, you may need to promote a slave to be the new master if the current master becomes unavailable. For example, this situation can occur when the original master cannot be restarted immediately due to a hardware problem, or when your application requires a feature that is only accessible via the master, such as the resource allocation facility or the ability to modify the deployment data.

To promote a slave to become the new master, you must shut down the slave and change its `IceGrid.Registry.ReplicaName` property to `Master` (or remove the property altogether). On restart, the new master notifies the nodes and registries that were active before it was shut down. An inactive registry or node will eventually connect to the new master if its default locator proxy contains the endpoint of the new master registry or the endpoint of a slave that is connected to the new master. If you cannot afford any down-time of the registry and want to minimize the down-time of the master, you should run at least two slaves. That way, if the master becomes unavailable, there will always be one registry available while you promote one of the slaves.

With Ice versions 3.5.0 and prior, a slave synchronizes its database upon connecting to the new master, therefore it is imperative that you promote a slave whose database is valid and up-to-date. To verify that the promoted master database is up-to-date, you can start the new master with the `--readonly` command-line option. While this option is in force, the new master does not update its database, and slaves do not synchronize their databases. You can review the master database with the IceGrid administrative tools and, if the deployment looks correct, you can restart the master without the `--readonly` option to permit updates and slave synchronization.

Starting with Ice 3.5.1, a check has been added to ensure that an out-of-date master database cannot overwrite a newer slave database. In such a scenario, an error message is printed out on the consoles of the IceGrid master and slave, and the slave connection to the master is aborted. You can still connect to the slave and master with the IceGrid administrative tools to verify their databases. Once you determine which database to keep, you can restart either the slave or the master with the `--initdb-from-replica=<name>` option to initialize the out-of-date database with the correct version. This check only works when both the master and slave use Ice 3.5.1 or later.

Note that there is nothing to prevent you from running two masters. If you start two masters and they contain different versions of the deployment information, some slaves and nodes might get updated with out-of-date deployment information (causing some of your servers to be deactivated). You can correct the problem by shutting down the faulty master, but it is important to keep this issue in mind when you restart a master since it might disrupt your applications.

See Also

- Registry Replication
- Resource Allocation using IceGrid Sessions
- Using IceGrid Deployment
- icegridadmin Command Line Tool

Copyright 2016, ZeroC, Inc.

IceGrid and the Administrative Facility

The Ice administrative facility provides a general purpose solution for administering individual Ice programs. IceGrid extends this functionality in several convenient ways:

- IceGrid automatically enables the facility in deployed servers, in its nodes and in its registry replicas.
- IceGrid uses the `Process` facet to terminate an active server, giving it an opportunity to perform an orderly shutdown.
- IceGrid provides a secure mechanism for invoking administrative operations on deployed servers, IceGrid nodes and IceGrid registry replicas.

The IceGrid administrative tools in turn use IceGrid's extended administrative facility to:

- display the properties of servers and services.
- attach remote loggers to the Logger of servers, services, nodes and registry replicas.
- manipulate and monitor IceBox services.

We discuss each of these items in separate sections below.

On this page:

- Enabling the Administrative Facility for a Deployed Server, Node or Registry Replica
  - Ice.Admin.Endpoints for Servers
  - admin Objects in Nodes and Registry Replicas
- Deactivating a Deployed Server
- Routing Administrative Requests
  - Obtaining a Proxy
  - Callbacks without Glacier2
  - Callbacks with Glacier2
- Using the Administrative Facility in IceGrid Utilities
  - Properties
  - Administering IceBox Services

Enabling the Administrative Facility for a Deployed Server, Node or Registry Replica

As we saw in our deployment example, the configuration properties for a deployed server include definitions for the following properties:

- `Ice.Admin.Endpoints`
- `Ice.Admin.ServerId`

The definition of `Ice.Admin.Endpoints` enables the Administrative Facility.

Ice.Admin.Endpoints for Servers

If a server's descriptor does not set `Ice.Admin.Enabled` and does not supply a value for `Ice.Admin.Endpoints`, IceGrid supplies a default value for `Ice.Admin.Endpoints` as shown below:

```
Ice.Admin.Endpoints=tcp -h 127.0.0.1
```

For security reasons, IceGrid specifies the local host interface (`127.0.0.1`) so that administrative access is limited to clients running on the same host. This configuration permits the IceGrid node to invoke operations on the server's `admin` object, but prevents remote access unless the client establishes an IceGrid administrative session.

Specifying a fixed port is unnecessary because the server registers its endpoints with IceGrid upon each new activation.

admin Objects in Nodes and Registry Replicas

Each IceGrid node and IceGrid registry replica provides by default an admin object hosted in the `IceGrid.Node` object adapter (for nodes) or in the `IceGrid.Registry.Internal` object adapter (for registry replicas). If you don't want to an admin object in a node or registry replica, set Ice.Admin.Enabled to 0 or a negative value in this node or registry's configuration.

Each of these admin objects carries all the built-in facets, currently `Logger`, Metrics, Process and Properties. Proxies to these admin objects can be retrieved through the `getNodeAdmin` and `getRegistryAdmin` operations described below.

Copyright 2016, ZeroC, Inc.

Deactivating a Deployed Server

An IceGrid node uses the `Ice::Process` interface to gracefully deactivate a server. In programs using Ice 3.3 or later, this interface is implemented by the administrative facet named `Process`. In earlier versions of Ice, an object adapter implemented this interface in a special servant if the adapter's `RegisterProcess` property was enabled.

Regardless of version, the Ice run time registers an `Ice::Process` proxy with the IceGrid registry when properly configured. Registration normally occurs during communicator initialization, but it can be delayed when a server needs to install its own administrative facets.

When the node is ready to deactivate a server, it invokes the `shutdown` operation on the server's `Ice::Process` proxy. If the server does not terminate in a timely manner, the node asks the operating system to terminate the process. Each server can be configured with its own deactivation timeout. If no timeout is configured, the node uses the value of the property `IceGrid.Node.WaitTime`, which defaults to `60` seconds.

If a server does not register an `Ice::Process` proxy, the IceGrid node cannot request a graceful termination and must resort instead to a more drastic, and potentially harmful, alternative by asking the operating system to terminate the server's process. On Unix, the node sends the `SIGTERM` signal to the process and, if the server does not terminate within the deactivation timeout period, sends the `SIGKILL` signal.

On Windows, the node first sends a `Ctrl+Break` event to the server and, if the server does not stop within the deactivation timeout period, terminates the process immediately.

Servers that disable the `Process` facet can install a signal handler in order to intercept the node's notification about pending deactivation. For example, portable C++ programs could use the `IceUtil::CtrlCHandler` class for this purpose. However, we recommend that servers be allowed to use the `Process` facet when possible.

Routing Administrative Requests

IceGrid defaults to using the local host interface when defining the endpoints of a deployed server's administrative object adapter. This configuration allows local clients such as the IceGrid node to access the server's `admin` object while preventing direct invocations from remote clients. A server's `admin` object may still be accessed remotely, but only by clients that establish an IceGrid administrative session. To facilitate these requests, IceGrid uses an intermediary object that relays requests to the server via its node. For example, the following figure illustrates the path of a `getProperty` invocation:

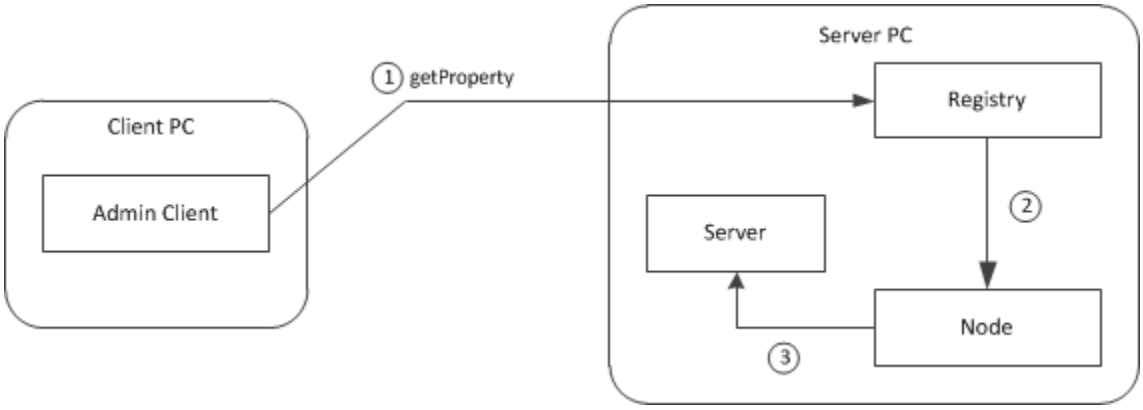

Routing for administrative requests on a server.

Obtaining a Proxy

During an administrative session, a client has two ways of obtaining the intermediary proxy for a server's `admin` object:

Copyright 2016, ZeroC, Inc.

Slice

```
module IceGrid {
    interface Admin {
        idempotent string getServerAdminCategory();
        idempotent Object* getServerAdmin(string id)
            throws ServerNotExistException, NodeUnreachableException,
DeploymentException;
        // ...
    };
};
```

If the client wishes to construct the proxy itself and already knows the server's ID, the client need only modify the proxy of the `IceGrid::Admin` object with a new identity. The identity's category must be the return value of `getServerAdminCategory`, while its name is the ID of the desired server. The example below demonstrates how to create the proxy and access the `Properties` facet of a server:

C++

```
IceGrid::AdminSessionPrx session = ...;
IceGrid::AdminPrx admin = session->getAdmin();
Ice::Identity serverAdminId;
serverAdminId.category = admin->getServerAdminCategory();
serverAdminId.name = "MyServerId";
Ice::PropertiesAdminPrx props =
    Ice::PropertiesAdminPrx::checkedCast(
        admin->ice_identity(serverAdminId), "Properties");
```

Alternatively, the `getServerAdmin` operation returns a proxy that refers to the `admin` object of the given server. This operation performs additional validation and therefore may raise one of the exceptions shown in its signature above.

A client can also obtain an intermediary proxy to the admin object of an IceGrid node or IceGrid registry replica by calling `getNodeAdmin` or `getRegistryAdmin`:

Slice

```
module IceGrid {
    interface Admin {
        idempotent Object* getNodeAdmin(string name)
            throws NodeNotExistException, NodeUnreachableException;

        idempotent Object* getRegistryAdmin(string name)
            throws RegistryNotExistException;
        // ...
    };
};
```

The name parameter corresponds to the node name or registry replica name. There is no operation comparable to `getServerAdminCateg`

Copyright 2016, ZeroC, Inc.

ory for IceGrid nodes and IceGrid registry replicas.

Callbacks without Glacier2

IceGrid also supports the relaying of callback requests from a back-end server to an administrative client over the client's existing connection to the registry, which is especially important for a client using a network port that is forwarded by a firewall or protected by a secure tunnel.

For this mechanism to work properly, a client that established its administrative session directly with IceGrid and not via a Glacier2 router mu st take additional steps to ensure that the proxies for its callback objects contain the proper identities and endpoints. The `IceGrid::Admin Session` interface provides an operation to help with the client's preparations:

Slice

```
module IceGrid {
interface AdminSession ... {
    idempotent Object* getAdminCallbackTemplate();
    // ...
};
};
```

As its name implies, the `getAdminCallbackTemplate` operation returns a *template proxy* that supplies the identity and endpoints a client needs to configure its callback objects. The information contained in the template proxy is valid for the lifetime of the administrative session. This operation returns a null proxy if the client's administrative session was established via a Glacier2 router, in which case the client should use the callback strategy described in the next section instead.

The endpoints contained in the template proxy are those of an object adapter in the IceGrid registry. The client must transfer these endpoints to the proxies for its callback objects so that callback requests from a server are sent first to IceGrid and then relayed over a bidirectional connection to the client, as shown below:

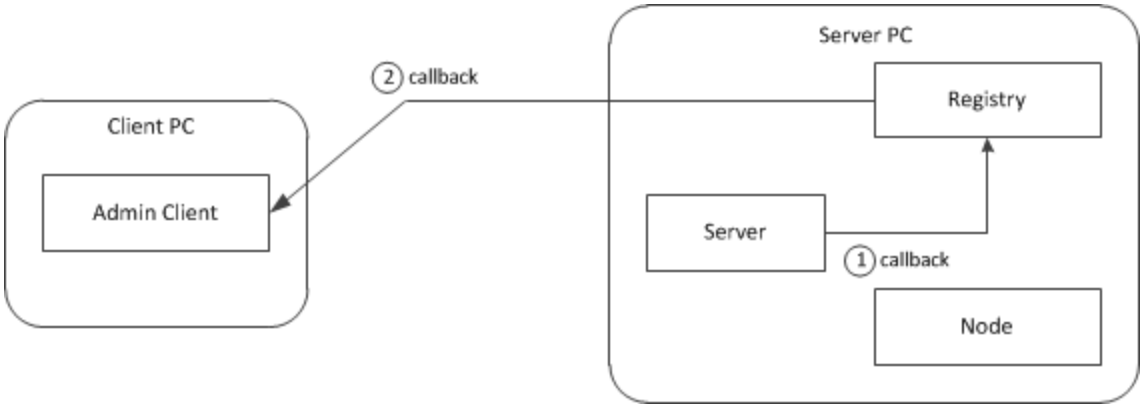

Routing for callback requests from a server.

Here is the complete list of steps:

1. Invoke `getAdminCallbackTemplate` to obtain the template proxy.
2. Extract the category from the template proxy's identity and use it in all callback objects.
3. Extract the endpoints from the template proxy and use them to establish the published endpoints of the callback object adapter.
4. Create the callback object adapter and associate it with the administrative session's connection, thereby establishing a bidirectional connection with IceGrid.
5. Add servants to the callback object adapter.

As an example, let us assume that we have deployed an IceBox server with the server id `icebox1` and our objective is to register a `Servic eObserver` callback that monitors the state of the IceBox services. The first step is to obtain a proxy for the administrative facet named `Ice Box.ServiceManager`:

Copyright 2016, ZeroC, Inc.

C++

```
IceGrid::AdminSessionPrx session = ...;
IceGrid::AdminPrx admin = session->getAdmin();
Ice::ObjectPrx obj = admin->getServerAdmin("icebox1");
IceBox::ServiceManagerPrx svcmgr =
    IceBox::ServiceManagerPrx::checkedCast(
        obj, "IceBox.ServiceManager");
```

Next, we retrieve the template proxy and compose the published endpoints for our callback object adapter:

C++

```
Ice::ObjectPrx tmpl = admin->getAdminCallbackTemplate();
Ice::EndpointSeq endpts = tmpl->ice_getEndpoints();
string publishedEndpoints;
for (Ice::EndpointSeq::const_iterator p = endpts.begin(); p != endpts.end(); ++p) {
    if (p == endpts.begin())
        publishedEndpoints = (*p)->toString();
    else
        publishedEndpoints += ":" + (*p)->toString();
}
communicator->getProperties()->setProperty(
    "CallbackAdapter.PublishedEndpoints", publishedEndpoints);
```

The final steps involve creating the callback object adapter, adding a servant, establishing the bidirectional connection and registering our callback with the service manager:

C++

```
Ice::ObjectAdapterPtr callbackAdapter =
communicator->createObjectAdapter("CallbackAdapter");
Ice::Identity cbid;
cbid.category = tmpl->ice_getIdentity().category;
cbid.name = "observer";
IceBox::ServiceObserverPtr obs = new ObserverI;
Ice::ObjectPrx cbobj = callbackAdapter->add(obs, cbid);
IceBox::ServiceObserverPrx cb = IceBox::ServiceObserverPrx::uncheckedCast(cbobj);
callbackAdapter->activate();
session->ice_getConnection()->setAdapter(callbackAdapter);
svcmgr->addObserver(cb);
```

At this point the client is ready to receive callbacks from the IceBox server whenever one of its services changes state.

Callbacks with Glacier2

A client that creates an administrative session via a Glacier2 router already has a bidirectional connection over which callbacks from

Copyright 2016, ZeroC, Inc.

administrative facets are relayed. The flow of requests is shown in the illustration below, which presents a simplified view with the router and IceGrid services all running on the same host.

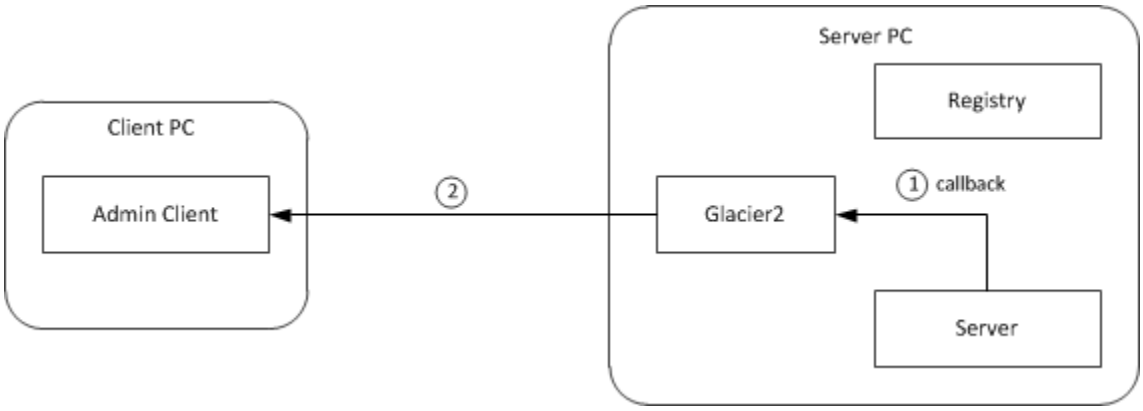

Routing for callback requests from a server.

To prepare for receiving callbacks, the client must perform the same steps as for any router client:

1. Obtain a proxy for the router.
2. Retrieve the category to be used in callback objects.
3. Create the callback object adapter and associate it with the router, thereby establishing a bidirectional connection.
4. Add servants to the callback object adapter.

Repeating the example from the previous section, we assume that we have deployed an IceBox server with the server ID `icebox1` and our objective is to register a `ServiceObserver` callback that monitors the state of the IceBox services. The first step is to obtain a proxy for the administrative facet named `IceBox.ServiceManager`:

C++

```
IceGrid::AdminSessionPrx session = ...;
IceGrid::AdminPrx admin = session->getAdmin();
Ice::ObjectPrx obj = admin->getServerAdmin("icebox1");
IceBox::ServiceManagerPrx svcmgr =
    IceBox::ServiceManagerPrx::checkedCast(obj, "IceBox.ServiceManager");
```

Now we are ready to create the object adapter and register the observer:

C++

```
Ice::RouterPrx router = communicator->getDefaultRouter();
Ice::ObjectAdapterPtr callbackAdapter =
    communicator->createObjectAdapterWithRouter("CallbackAdapter", router);
Ice::Identity cbid;
cbid.category = router->getCategoryForClient();
cbid.name = "observer";
IceBox::ServiceObserverPtr obs = new ObserverI;
Ice::ObjectPrx cbobj = callbackAdapter->add(obs, cbid);
IceBox::ServiceObserverPrx cb = IceBox::ServiceObserverPrx::uncheckedCast(cbobj);
callbackAdapter->activate();
svcmgr->addObserver(cb);
```


Copyright 2016, ZeroC, Inc.

At this point the client is ready to receive callbacks from the IceBox server whenever one of its services changes state.

Using the Administrative Facility in IceGrid Utilities

This section discusses the ways in which the IceGrid utilities make use of the administrative facility.

Properties

The command line and graphical utilities allow you to explore the configuration properties of a server or service.

One property in particular, `BuildId`, is given special consideration by the graphical utility. Although it is not used by the Ice run time, the `BuildId` property gives you the ability to describe the build configuration of your application. The property's value is shown by the graphical utility in its own field in the attributes of a server or service, as well as in the list of properties. You can also retrieve the value of this property using the command-line utility with the following statement:

```
> server property MyServerId BuildId
```

Or, for an IceBox service, with this command:

```
> service property MyServerId MyService BuildId
```

The utilities use the `Properties` facet to access these properties, via a proxy obtained as described above.

Administering IceBox Services

IceBox provides an administrative facet that implements the `IceBox::ServiceManager` interface, which supports operations for stopping an active service, and for starting a service that is currently inactive. These operations are available in both the command line and graphical utilities.

IceBox also defines a `ServiceObserver` interface for receiving callbacks when services are stopped or started. The graphical utility implements this interface so that it can present an updated view of the state of an IceBox server. We presented examples that demonstrate how to register an observer with the IceBox administrative facet.

See Also

- Administrative Facility
- The Process Facet
- The Properties Facet
- Creating the admin Object
- The admin Object
- Custom Administrative Facets
- Security Considerations for Administrative Facets
- Portable Signal Handling in C++
- Bidirectional Connections
- Using IceGrid Deployment
- Glacier2 Integration with IceGrid
- IceGrid Administrative Sessions
- icegridadmin Command Line Tool
- Callbacks through Glacier2
- IceBox
- IceBox Administration
- IceGrid.*

Copyright 2016, ZeroC, Inc.

Securing IceGrid

IceGrid's registry and node services expose multiple network endpoints that a malicious client could use to gain access to IceGrid functionality and interfere with deployed applications. This presents a significant security risk in network environments that are exposed to untrusted clients. For example, a malicious client could connect to a node and use IceGrid's internal interfaces to deploy and run its own server executable.

This page describes the steps you can take to secure your IceGrid application.

On this page:

- IceGrid Security Overview
- Understanding the Registry Endpoints
  - Client Endpoint
  - Server Endpoint
  - Internal Endpoint
  - Session Manager Endpoint
  - Administrative Session Manager Endpoint
  - IceLocatorDiscovery Endpoint
  - Outgoing Connections
- Understanding the Node Endpoints
- Understanding the Administrative Endpoints with IceGrid

IceGrid Security Overview

Using a firewall is one way to prevent unauthorized use of IceGrid's facilities. Another solution is to use IceSSL: you can generate SSL certificates for each component and configure them to trust and accept connections only from other authorized components. The remainder of this section discusses the IceSSL solution but also provides useful information for those interested in securing IceGrid with a firewall.

To restrict access using IceSSL, we need to establish trust relationships between IceGrid registry replicas, nodes, and deployed servers. IceSSL allows us to do this using configuration properties. The trust relationships are based on the information contained in SSL certificates.

There are several possible strategies for generating certificates. At a minimum you will need the following:

- one certificate for all of the registries
- one certificate for all of the nodes
- one certificate for all of the servers managed by IceGrid

The certificates that you generate for registries and nodes should be protected with a password to ensure that only privileged users can start these services. However, we do not recommend using a password to protect the certificate for deployed servers because it would need to be specified in clear text in each server's configuration (servers that are activated by IceGrid must not prompt for a password). Furthermore, this password might appear in multiple places, such as an XML descriptor file, the IceGrid registry database, and property files generated by IceGrid nodes. The complexity involved in protecting access to every file that contains a clear text password could be overwhelming. Instead, we recommend that you protect access to the server certificate using file system permissions.

Depending on your organization and the roles of each person that uses IceGrid, you may decide to create additional certificates. For example, you might create a unique certificate for each IceGrid node instance if you deploy nodes on end-user machines and wish to configure the IceGrid registry to authorize connections only from the nodes of trusted users.

You can use the `iceca` script to establish a certificate authority and generate certificates. The sections that follow describe the interactions between the registry, node, and servers, and show how to configure IceSSL to restrict access to trusted peers. For the purposes of this discussion, we assume that the SSL certificates use the common names shown below:

- `IceGrid Registry`
- `IceGrid Node`
- `Server`

The Ice distribution includes a C++ example that demonstrates how to configure a secure IceGrid deployment in the `demo/IceGrid/secure` subdirectory. This example includes a script to generate certificates for a registry, a node, a Glacier2 router, and a server. For more information, see the `README` file provided with the example.

Understanding the Registry Endpoints

The IceGrid registry has three mandatory endpoints representing the client, server, and internal endpoints. The registry also has two optional endpoints (the session manager and administrative session manager endpoints) that are only useful when accessing IceGrid via Glacier2.

Client Endpoint

Copyright 2016, ZeroC, Inc.

The registry client endpoint is used by Ice applications that create client sessions in order to use the resource allocation facility. It is also used by administrative clients that create sessions for managing the registry. Finally, the client endpoint is used by Ice applications that use the `IceGrid::Query` interface or resolve indirect proxies via the IceGrid locator.

Two distinct permission verifiers authorize the creation of client sessions and administrative sessions. The remaining functionality available via the client endpoint, such as resolving objects and object adapters using the `IceGrid::Query` interface or the Ice locator mechanism, is accessible to any client that is able to connect to the client endpoint.

It is safe to use an insecure transport for the client endpoint if it is only being used for locator queries. However, you should use a secure transport if you have enabled client and administrative sessions (by configuring the appropriate permission verifiers). Creating a session over an insecure transport poses a security risk because the user name and password are sent in clear text over the network.

If you include secure and insecure transports in the registry's client endpoints, you should ensure that applications that need to authenticate with IceGrid permission verifiers use a secure transport.

It is not necessary to restrict SSL access to the client endpoints (using the property `IceSSL.TrustOnly.Server.IceGrid.Registry.Client`) as long as you use client and administrative permission verifiers for authentication. This property is only useful for restricting access to client and administrative sessions when using null permission verifiers. Note however that if both client and administrative sessions are enabled, you will only be able to restrict access to one set of clients since you cannot distinguish clients that create client sessions from clients that create administrative sessions.

Server Endpoint

Ice servers use the registry's server endpoint to register their object adapter endpoints and send information to administrative clients connected via the registry.

Securing this endpoint with IceSSL is necessary to prevent a malicious program from potentially hijacking a server by registering its endpoints first. The property definition shown below demonstrates how to limit access to this endpoint to trusted Ice servers:

```
IceSSL.TrustOnly.Server.IceGrid.Registry.Server=CN="Server"
```

Internal Endpoint

IceGrid nodes and registry replicas use the internal endpoint to communicate with the registry. For example, nodes connect to the internal endpoint of each active registry, and registry slaves establish a session with their master via this endpoint.

The internal endpoint must be secured with IceSSL to prevent malicious Ice applications from gaining access to sensitive functionality that is intended to be used only by nodes and registry replicas. You can restrict access to this endpoint with the following property:

```
IceSSL.TrustOnly.Server.IceGrid.Registry.Internal=CN="IceGrid Node";CN="IceGrid
Registry"
```

Session Manager Endpoint

The session manager endpoint is used by Glacier2 to create IceGrid client sessions. The functionality exposed by this endpoint is unrestricted so you must either secure it or disable it (this endpoint is disabled by default). The property shown below demonstrates how to configure IceSSL so that only Glacier2 routers are accepted by this endpoint:

```
IceSSL.TrustOnly.Server.IceGrid.Registry.SessionManager=CN="Glacier2 Router Client"
```

In this example, `Glacier2 Router Client` is the common name of the Glacier2 router used by clients to create IceGrid client sessions.

Administrative Session Manager Endpoint

Glacier2 routers use the registry's administrative session manager endpoint to create IceGrid administrative sessions. The functionality

Copyright 2016, ZeroC, Inc.

exposed by this endpoint is unrestricted, so you must either secure it or disable it (this endpoint is disabled by default). The property shown below demonstrates how to configure IceSSL so that only Glacier2 routers are accepted by this endpoint:

```
IceSSL.TrustOnly.Server.IceGrid.Registry.AdminSessionManager=CN="Glacier2 Router
Admin"
```

In this example, `Glacier2 Router Admin` is the common name of the Glacier2 router used by clients to create IceGrid administrative sessions. Note that if you use a single Glacier2 router instance for both client and administrative sessions, you will need to use the same common name to restrict access to both session manager endpoints:

```
IceSSL.TrustOnly.Server.IceGrid.Registry.SessionManager=CN="Glacier2 Router Client"
IceSSL.TrustOnly.Server.IceGrid.Registry.AdminSessionManager=CN="Glacier2 Router
Client"
```

IceLocatorDiscovery Endpoint

The registry (including all replicas) listens by default for UDP multicast requests from IceLocatorDiscovery clients. Each client request includes a proxy to which the registry sends its response; by default, the client's "response callback" proxy uses a UDP unicast endpoint but it can be configured to use a different transport. The registry invokes on the client's response callback proxy and provides a proxy of its own containing the registry's client endpoint.

Ice does not support a secure multicast transport therefore these discovery requests cannot be encrypted or restricted via trust relationships. You can prevent registries from listening for discovery requests by setting `IceGrid.Registry.Discovery.Enabled`:

```
IceGrid.Registry.Discovery.Enabled=0
```

Outgoing Connections

The registry establishes outgoing connections to other registries and nodes. You should configure the `IceSSL.TrustOnly.Client` property to restrict connections to these trusted peers:

```
IceSSL.TrustOnly.Client=CN="IceGrid Registry";CN="IceGrid Node"
```

The registry can also connect to Glacier2 routers and permission verifier objects. To allow connections to these services, you must include in this property the common names of Glacier2 routers that create client or administrative sessions, as well as the common names of servers that host the permission verifier objects.

Understanding the Node Endpoints

An IceGrid node has only one endpoint, which is used for internal communications with the registry. As a result, it should be configured to accept connections only from IceGrid registries:

```
IceSSL.TrustOnly.Server=CN="IceGrid Registry"
```

A node also establishes outgoing connections to the registry's internal endpoint, as well as the `Ice.Admin` endpoint of deployed servers.

Copyright 2016, ZeroC, Inc.

You should configure the `IceSSL.TrustOnly.Client` property as shown below to verify the identity of these peers:

```
IceSSL.TrustOnly.Client=CN="Server";CN="IceGrid Registry"
```

Understanding the Administrative Endpoints with IceGrid

By default, IceGrid sets the endpoints of a deployed server's `Ice.Admin` adapter to `tcp -h 127.0.0.1`. This setting is already quite secure because it only accepts connections from processes running on the same host. However, since you already need to configure IceSSL so that a server can authenticate with the IceGrid registry (servers connect to the registry to register their endpoints), you might as well use a secure endpoint for the `Ice.Admin` adapter and configure it to accept connections only from IceGrid nodes:

```
IceSSL.TrustOnly.Server.Ice.Admin=CN="IceGrid Node"
```

This is only necessary if the `Ice.Admin` endpoint is enabled (which it is by default).

You can also set the `IceSSL.TrustOnly.Client` property so that the server is only permitted to connect to the IceGrid registry:

```
IceSSL.TrustOnly.Client=CN="IceGrid Registry"
```

If your server invokes on other servers, you will need to modify this setting to allow secure connections to them.

See Also

- IceSSL
- Configuring IceSSL
- Setting up a Certificate Authority
- Glacier2 Integration with IceGrid
- Resource Allocation using IceGrid Sessions
- IceGrid Administrative Sessions
- Well-Known Objects
- IceGrid and the Administrative Facility
- Registry Replication
- IceSSL.*

Copyright 2016, ZeroC, Inc.

icegridadmin Command Line Tool

The `icegridadmin` utility is a command-line tool for administering an IceGrid domain. Deploying an application with this utility requires an XML file that defines the descriptors.

On this page:

- Usage
- Application Commands
- Node Commands
- Registry Commands
- Server Commands
- Service Commands
- Adapter Commands
- Object Commands
- Server Template
- Service Template

Usage

The IceGrid administration tool supports the following command-line options:

```
Usage: icegridadmin [options]
Options:
-h, --help           Show this message.
-v, --version        Display the Ice version.
-e COMMANDS          Execute COMMANDS.
-d, --debug          Print debug messages.
-s, --server         Start icegridadmin as a server (to parse XML files).
-i, --instanceName   Connect to the registry with the given instance name.
-H, --host           Connect to the registry at the given host.
-P, --port           Connect to the registry running on the given port.
-u, --username       Login with the given username.
-p, --password       Login with the given password.
-S, --ssl            Authenticate through SSL.
-r, --replica NAME   Connect to the replica NAME.
```

The `-e` option causes the tool to execute the given commands and then exit without entering an interactive mode. The `-s` option starts `icegridadmin` in a server mode that supports the `IceGrid::FileParser` interface; a proxy for the object is printed to standard output. If neither `-e` nor `-s` is specified, the tool enters an interactive mode in which you issue commands at a prompt.

To communicate with the IceGrid registry, `icegridadmin` establishes an administrative session. The tool uses SSL authentication if you specify the `-S` option or define its equivalent property `IceGridAdmin.AuthenticateUsingSSL`. Otherwise, `icegridadmin` uses password authentication and prompts you for the username and password if you do not specify them via command-line options or properties. If you want `icegridadmin` to establish its session using a Glacier2 router, define `Ice.Default.Router` appropriately.

If specified, the following command-line options override their property equivalents:

Option	Property
`-i, --instanceName`	`IceGridAdmin.InstanceName`
`-H, --host`	`IceGridAdmin.Host`
`-P, --port`	`IceGridAdmin.Port`
`-u, --username`	`IceGridAdmin.Username`
`-p, --password`	`IceGridAdmin.Password`

Copyright 2016, ZeroC, Inc.

`-S, --ssl`	`IceGridAdmin.AuthenticateUsingSSL`
`-r, --replica`	`IceGridAdmin.Replica`

`icegridadmin` determines its target registry as follows:

1. Connect via a Glacier2 router if `Ice.Default.Router` is defined.
2. Otherwise, if `Ice.Default.Locator` is defined, connect to the specified registry.
3. Otherwise, if a host is defined via the `-H` or `--host` options or the equivalent property, connect to the registry at the specified host. If no port is defined via the `-P` or `--port` options or the equivalent property, `icegridadmin` uses the standard IceGrid TCP (4061) or SSL (4062) port.
4. Otherwise, `icegridadmin` attempts to locate a registry by issuing a UDP multicast discovery request. (`icegridadmin` does not use the IceGridDiscovery plug-in.) If the tool discovers more than one registry, it presents a list and asks you to select one.

See IceGridAdmin.* for more information on the configuration properties supported by `icegridadmin`.

Once it has connected to the registry and successfully established a session, `icegridadmin` displays its command prompt. The `help` command displays the following usage information:

- `help`
  Print this message.
- `exit`, `quit`
  Exit this program.
- `CATEGORY help`
  Print the help section of the given `CATEGORY`
- `COMMAND help`
  Print the help of the given `COMMAND`.

The tool's commands are organized by category. The supported command categories are shown below:

- `application`
- `node`
- `registry`
- `server`
- `service`
- `adapter`
- `object`
- `server template`
- `service template`

You can obtain more information about each category using the `help` command:

```
>>> application help
```

Application Commands

- `application add [-n | --no-patch] DESC [TARGET ... ] [NAME=VALUE ... ]`
  Add applications described in the XML descriptor file `DESC`. If specified the optional targets are deployed. Variables are defined using the `NAME=VALUE` syntax. The application is automatically patched unless the `-n` or `--no-patch` option is used to disable it.
- `application remove NAME`
  Remove the application named `NAME`.
- `application describe NAME`
  Describe the application named `NAME`.
- `application diff [-s | --servers] DESC [TARGET ...] [NAME=VALUE ...]`
  Print the differences between the application in the XML descriptor file `DESC` and the current deployment. If `-s` or `--servers` is specified, print the the list of servers affected by the differences. Variables are defined using the `NAME=VALUE` syntax.

Copyright 2016, ZeroC, Inc.

- `application update [-n | --no-restart] DESC [TARGET ...] [NAME=VALUE ...]`
  Update the application in the XML descriptor file `DESC`. If `-n` or `--no-restart` is specified, the update will fail if it would require restarting one or more servers. Variables are defined using the `NAME=VALUE` syntax. Use the `diff --servers` command to discover which servers would be affected by an update, including those that would require a restart.
- `application patch [-f | --force] NAME`
  Patch the application named `NAME`. If `-f` or `--force` is specified, IceGrid will first shut down any servers that depend on the data to be patched.
- `application list`
  List all deployed applications.

Node Commands

- `node list`
  List all registered nodes.
- `node describe NAME`
  Show information about node `NAME`.
- `node ping NAME`
  Ping node `NAME`.
- `node load NAME`
  Print the load of the node `NAME`.
- `node sockets [NAME]`
  Print the number of processor sockets for node `NAME`. If `NAME` is omitted, print the number of processor sockets for each node. (The `IceGrid.Node.ProcessorSocketCount` property allows you to explicitly set this value for systems where the number of sockets cannot be obtained programmatically.)
- `node show [OPTIONS] NAME [log | stderr | stdout]`
  Print the Ice log messages of the node (with `log`), or print the text from the node's standard error or standard output (with `stderr` or `stdout`). The supported options are shown below:
  - `-f`, `--follow`
    With `log`, create a remote logger that prints each new log message.
    With `stderr` or `stdout`, wait for new text to be available in the file where `stderr` or `stdout` is redirected.
  - `-t`, `--tail N`
    Print the last `N` log messages (for `log`) or N lines of text (for `stderr` or `stdout`)
  - `-h`, `--head N`
    Print the first `N` lines of text (invalid option with `log`).
- `node shutdown NAME`
  Shutdown node `NAME`.

Registry Commands

- `registry list`
  List all registered registries.
- `registry describe NAME`
  Show information about registry `NAME`.
- `registry ping NAME`
  Ping registry `NAME`.
- `registry show [OPTIONS] NAME [log | stderr | stdout]`
  Print the Ice log messages of the registry (with `log`), or print the text from the registry's standard error or standard output (with `stderr` or `stdout`). The supported options are shown below:
  - `-f`, `--follow`
    With `log`, create a remote logger that prints each new log message.
    With `stderr` or `stdout`, wait for new text to be available in the file where `stderr` or `stdout` is redirected.
  - `-t`, `--tail N`
    Print the last `N` log messages (for `log`) or N lines of text (for `stderr` or `stdout`)
  - `-h`, `--head N`
    Print the first `N` lines of text (invalid option with `log`).

Copyright 2016, ZeroC, Inc.

- `registry shutdown NAME`
  Shutdown registry `NAME`.

Server Commands

- `server list`
  List all registered servers.
- `server remove ID`
  Remove server `ID`.
- `server describe ID`
  Describe server `ID`.
- `server properties ID`
  Get the run-time properties of server `ID`.
- `server property ID NAME`
  Get the run-time property `NAME` of server `ID`.
- `server state ID`
  Get the state of server `ID`.
- `server pid ID`
  Get the process ID of server `ID`.
- `server start ID`
  Start server `ID`.
- `server stop ID`
  Stop server `ID`.
- `server patch ID`
  Patch server `ID`.
- `server signal ID SIGNAL`
  Send `SIGNAL` (such as `SIGTERM` or `15`) to server `ID`.
- `server stdout ID MESSAGE`
  Write `MESSAGE` on server `ID`'s standard output.
- `server stderr ID MESSAGE`
  Write `MESSAGE` on server `ID`'s standard error.
- `server show [OPTIONS] ID [log | stderr | stdout | LOGFILE]`
  Print the Ice log messages of the server (with `log`), or print the text from the server's standard error, standard output or the log file LOGFILE (with `stderr`, `stdout` or `LOGFILE`). The supported options are shown below:
  - `-f`, `--follow`
    With `log`, create a remote logger that prints each new log message.
    With `stderr`, `stdout` and `LOGFILE`, wait for new text to be available in the file.
  - `-t`, `--tail N`
    Print the last `N` log messages (for `log`) or N lines of text (for `stderr`, `stdout` or `LOGFILE`)
  - `-h`, `--head N`
    Print the first `N` lines of text (invalid option with `log`).
- `server enable ID`
  Enable server `ID`.
- `server disable ID`
  Disable server `ID`. A disabled server cannot be started on demand.

Service Commands

- `service start ID NAME`
  Starts service `NAME` in IceBox server `ID`.
- `service stop ID NAME`
  Stops service `NAME` in IceBox server `ID`.

Copyright 2016, ZeroC, Inc.

- `service describe ID NAME`
  Describes service `NAME` in IceBox server `ID`.
- `service properties ID NAME`
  Get the run-time properties of service `NAME` from IceBox server `ID`.
- `service property ID NAME PROPERTY`
  Get the run-time property `PROPERTY` of service `NAME` from IceBox server `ID`.
- `service list ID`
  List the services in IceBox server `ID`.

Adapter Commands

- `adapter list`
  List all registered adapters.
- `adapter endpoints ID`
  Show the endpoints of adapter or replica group `ID`.
- `adapter remove ID`
  Remove adapter or replica group `ID`.

Object Commands

The `object` command operates on well-known objects.

- `object add PROXY [TYPE]`
  Add a well-known object to the registry, optionally specifying its type.
- `object remove IDENTITY`
  Remove a well-known object from the registry.
- `object find TYPE`
  Find all well-known objects with the type `TYPE`.
- `object describe EXPR`
  Describe all well-known objects whose stringified identities match the expression `EXPR`. A trailing wildcard is supported in `EXPR`, for example "`object describe Ice*`".
- `object list EXPR`
  List all well-known objects whose stringified identities match the expression `EXPR`. A trailing wildcard is supported in `EXPR`, for example "`object list Ice*`".

Server Template

- `server template instantiate APPLICATION NODE TEMPLATE  [NAME=VALUE ...]`
  Instantiate the requested server template defined in the given application on a node. Variables are defined using the `NAME=VALUE` s yntax.
- `server template describe APPLICATION TEMPLATE`
  Describe a server template`TEMPLATE` from the given application.

Service Template

- `service template describe APPLICATION TEMPLATE`
  Describe a service template`TEMPLATE` from the given application.

See Also

- IceGrid Administrative Sessions
- Glacier2 Integration with IceGrid
- IceGrid XML Features
- Using Descriptor Variables and Parameters

Copyright 2016, ZeroC, Inc.

- Application Distribution
- Getting Started with IceGrid
- IceGridAdmin.*
- IceLocatorDiscovery

Copyright 2016, ZeroC, Inc.

IceGrid Admin Graphical Tool

IceGrid Admin is a graphical client for IceGrid. IceGrid Admin allows you to monitor servers deployed by IceGrid, create new IceGrid applications, deploy these applications, and more.

Topics

- Getting Started with IceGrid Admin
- Live Deployment Tab
  - Connection to an IceGrid Registry
  - Run Time Components
    - Registry Run Time Component
    - Slave Registry Run Time Component
    - Node Run Time Component
    - Server Run Time Component
    - Adapter Run Time Component
    - Database Environment Run Time Component
    - Service Run Time Component
    - Metrics View Run Time Component
  - Application Component
  - Ice Log Dialog
  - Log File Dialog
  - Metrics Graph
- Application Tabs
  - Editing and Saving IceGrid Descriptors
  - Navigation within an Application Tab
  - IceGrid Descriptors
    - Variables in IceGrid Descriptors
    - Application Descriptor
    - Node Descriptor
    - Server Descriptor
    - Service Descriptor
    - Adapter Descriptor
    - Database Environment Descriptor
    - Property Set Descriptor
    - Replica Group Descriptor
    - Server Template Descriptor
    - Service Template Descriptor

Copyright 2016, ZeroC, Inc.

Getting Started with IceGrid Admin

This page describes how to launch the IceGrid Admin tool.

On this page:

- System Requirements
- Starting IceGrid Admin
- Command Line Arguments
- Main IceGrid Admin Window
  - Tabs
  - Status Bar

System Requirements

IceGrid Admin is a Java™ application supported on a wide range of platforms, including Windows, Linux and OS X.

The minimum requirements for running IceGrid Admin are listed below:

- `icegridgui.jar`, usually installed in the `bin` or `lib` directory of your Ice installation
- Java SE Runtime Environment 7 or later

In order to use IceGrid Admin's metrics graphs feature, you will need the JavaFX Runtime Environment, bundled with recent updates of the Oracle Java SE Runtime Environment 7 on Windows, Linux and OS X.

You can download Oracle Java SE for most platforms from Oracle.

If you want to read IceGrid XML files from IceGrid Admin, you also need to have the `icegridadmin` command-line utility (version 3.6) in your PATH.

Starting IceGrid Admin

On Windows IceGrid Admin can be started by clicking the IceGrid Admin icon in the Start menu:

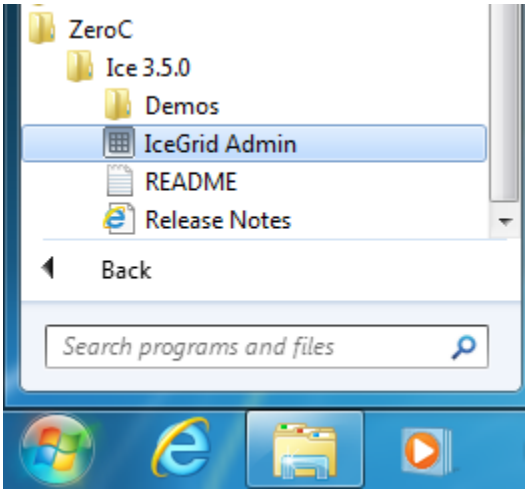

On OS X IceGrid Admin can be started by clicking the IceGrid Admin icon in Finder Applications folder:

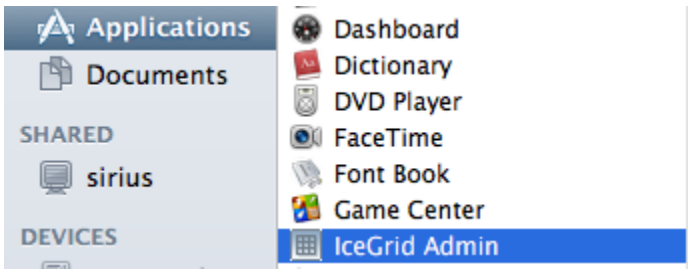

Copyright 2016, ZeroC, Inc.

On Linux with a RPM installation, you can use the `icegridgui` script installed in `/usr/bin`:

```
$ icegridgui
```

On all platforms, you can also start IceGrid Admin from a terminal by typing:

```
java -jar path-to-icegridgui.jar
```

Command Line Arguments

IceGrid Admin can be configured using Ice properties, and like with most Ice applications, these properties can be set using command-line arguments or a configuration file (or both).

Since IceGrid Admin is an Ice application (a client to the IceGrid registry), setting regular Ice properties can be useful as well. For example, you can set the `Ice.Trace.Network` property to get detailed information about network communications sent to `stderr`.

```
$ icegridgui --Ice.Trace.Network=2
```

There are also a number of properties specific to IceGrid Admin itself, describe in IceGridAdmin.*.

If you need to set many properties, it is a good idea to write a configuration file and use the `--Ice.Config` command-line argument to specify the location of this file. For example:

```
> java -jar "C:\Program Files (x86)\ZeroC\Ice-3.6b\bin\icegridgui.jar"
--Ice.Config=icegridadmin.conf
```

Main IceGrid Admin Window

The main IceGrid Admin window allows to navigate between your live deployment and the definitions of several applications.

Tabs

The main IceGrid Admin window shows one or more tabs:

Copyright 2016, ZeroC, Inc.

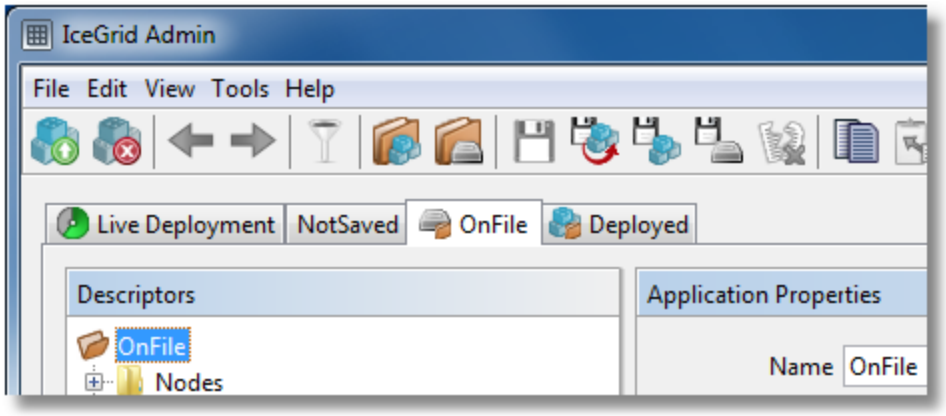

- The Live Deployment tab ( ) displays information about an IceGrid deployment you have logged into. There is always one and only one Live Deployment tab. When you are not connected to an IceGrid deployment, the corresponding pane is empty.
- A Live Application tab ( ) displays application definitions retrieved from the IceGrid registry you are connected to. As long as you do not change anything in the associated pane, IceGrid Admin will keep the information up-to-date. For example if another administrator adds a new server definition in this application definition, it will appear automatically and immediately in this pane.
- A File-Based Application tab ( ) displays application definitions retrieved from an IceGrid XML file.
- An icon-less tab displays the definitions of an application that is not bound to an IceGrid registry or to a file, such as a brand new application. A live application with unsaved modifications also becomes icon-less if the connection to its IceGrid registry is lost.

IceGrid Admin may show any number of application tabs, including none at all.

Status Bar

The status bar at the bottom of the main window shows information about operations performed by IceGrid Admin, or messages received from the IceGrid registry.

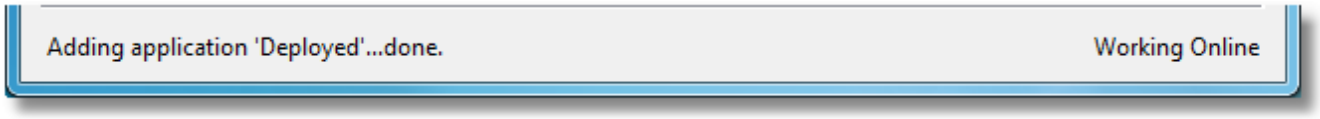

See Also

- IceGridAdmin.*
- icegridadmin Command Line Tool

Copyright 2016, ZeroC, Inc.

Live Deployment Tab

The Live Deployment tab shows the run-time status and configuration of an existing IceGrid deployment, and allows you to perform various administrative tasks on this deployment, such as:

- start or stop a server
- retrieve metrics associated with a server
- display metrics from one or more servers on a graph
- enable or disable a server
- send signal to a server
- retrieve the log files of a server
- patch an entire application
- register a new "well-known" object with the IceGrid registry
- undeploy a deployed application

Topics

- Connection to an IceGrid Registry
- Run Time Components
  - Registry Run Time Component
  - Slave Registry Run Time Component
  - Node Run Time Component
  - Server Run Time Component
  - Adapter Run Time Component
  - Database Environment Run Time Component
  - Service Run Time Component
  - Metrics View Run Time Component
- Application Component
- Ice Log Dialog
- Log File Dialog
- Metrics Graph

Copyright 2016, ZeroC, Inc.

Connection to an IceGrid Registry

In order to administer or monitor an IceGrid deployment, you need to connect to the IceGrid registry of this deployment. If your IceGrid registry is replicated, you can connect to any replica for monitoring purposes; if you intend to change definitions, for example to describe a new server, you need to connect to the master IceGrid registry. This page describes how to connect to an IceGrid registry.

On this page:

- Connecting using a Saved Connection
- Creating a new Connection
  - TCP Connection to a local IceGrid registry
  - SSL Connection to local IceGrid using X.509 Credentials
  - SSL Connection through Glacier2 router
- SSL Connections and Certificates
- Editing a Saved Connection
- Closing a Connection

Connecting using a Saved Connection

Use `File > Login...` or press the button to open the `Saved Connections` dialog:

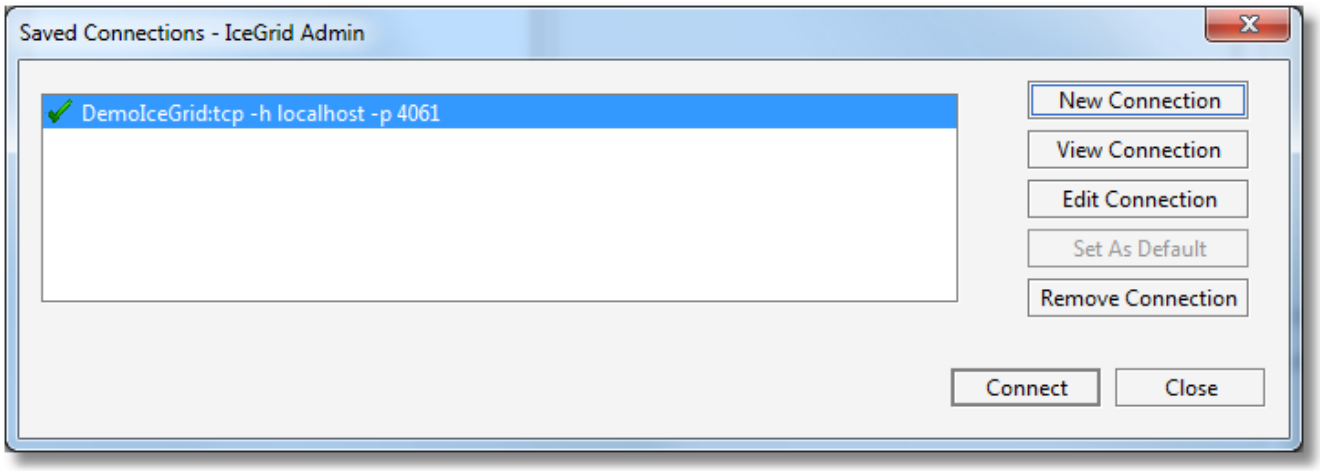

Double-click on the IceGrid registry you want to connect to, or select the IceGrid registry and click on the `Connect` button.

If the connection contains saved credentials, IceGrid Admin will immediately attempt to connect to the selected IceGrid registry with these credentials. Otherwise, it will open a new dialog to request the missing password (two passwords in some cases), such as:

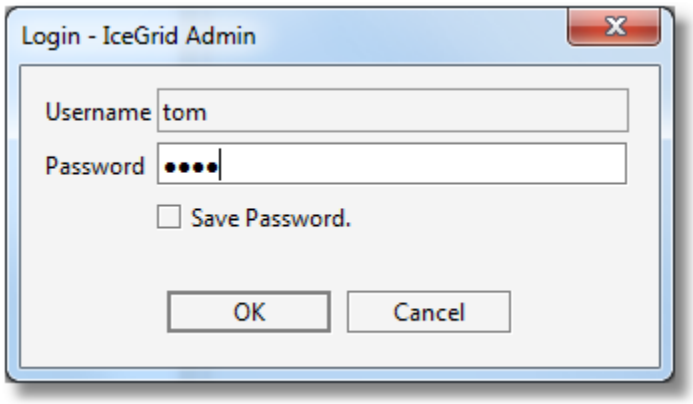

Creating a new Connection

Copyright 2016, ZeroC, Inc.

In the `Saved Connections` dialog, click on `New Connection` to open the New Connection wizard.

We describe three typical connection scenarios in the sections below.

TCP Connection to a local IceGrid registry

Let's create a TCP connection to the IceGrid registry running on localhost.

Step 1: Select Direct Connection:

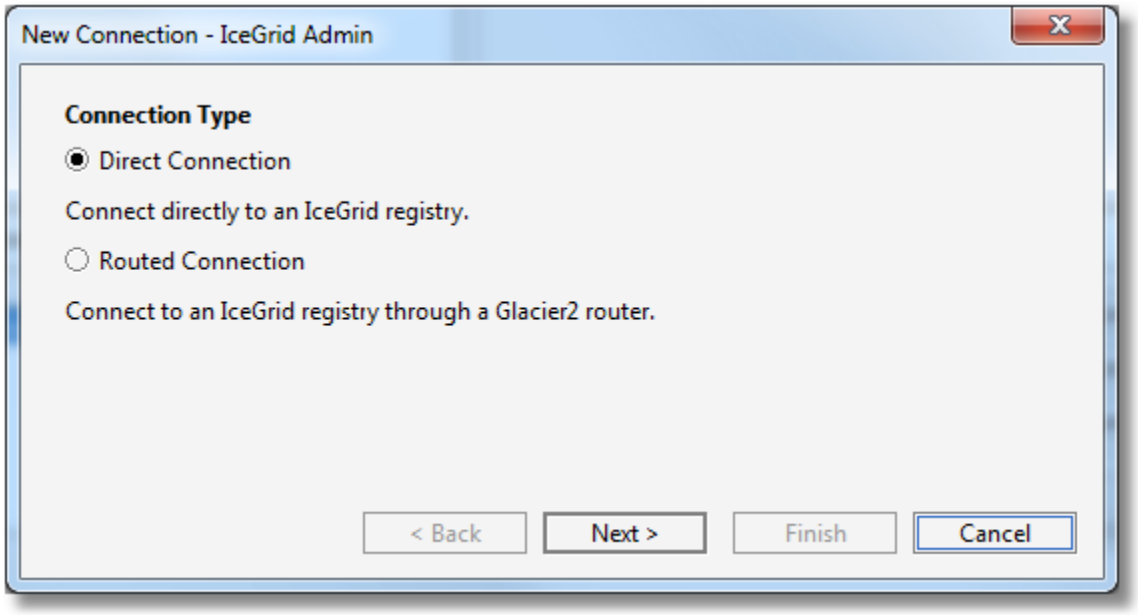

Step 2: Check the box to connect to the master registry:

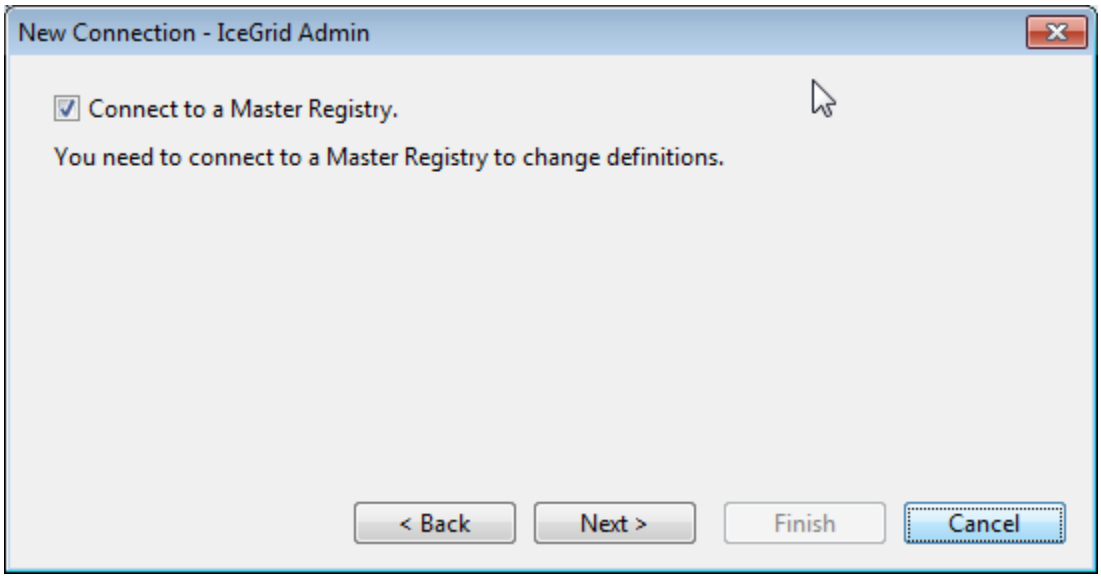

If you have several replicas this ensures you will always connect to the master registry.

Step 3: Select Manual Endpoint to manually enter address information:

Copyright 2016, ZeroC, Inc.

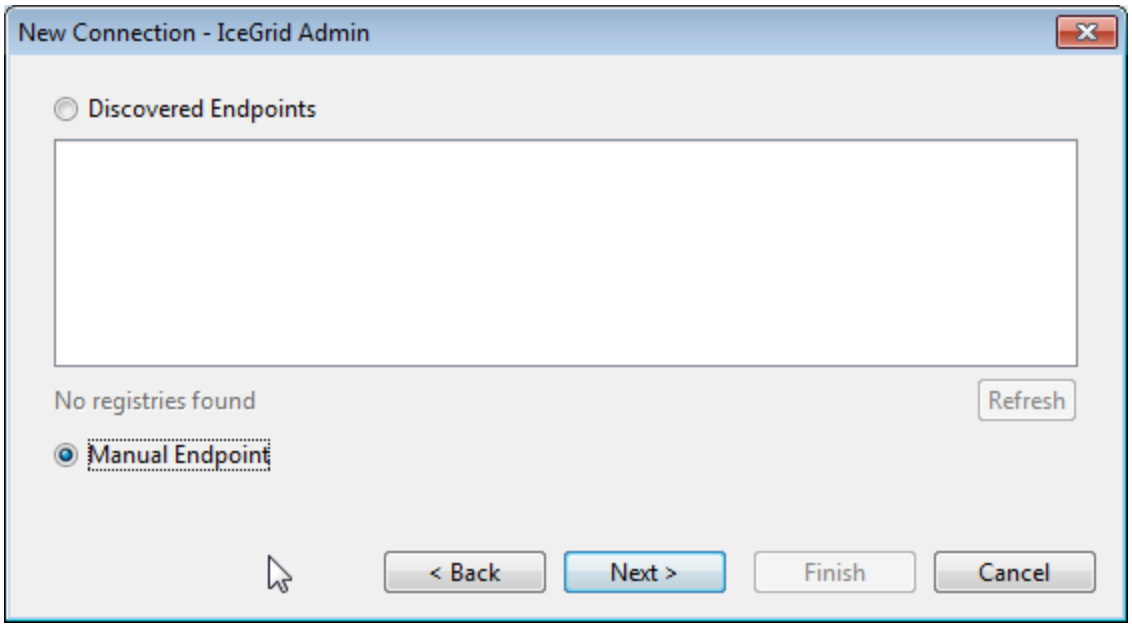

As of Ice 3.6, IceGrid Admin can also discover registries using UDP multicast. Any registries it finds will be shown in the list, or click the Refresh button to search again.

Step 4: Select the first option to enter the addressing information as a hostname and port number:

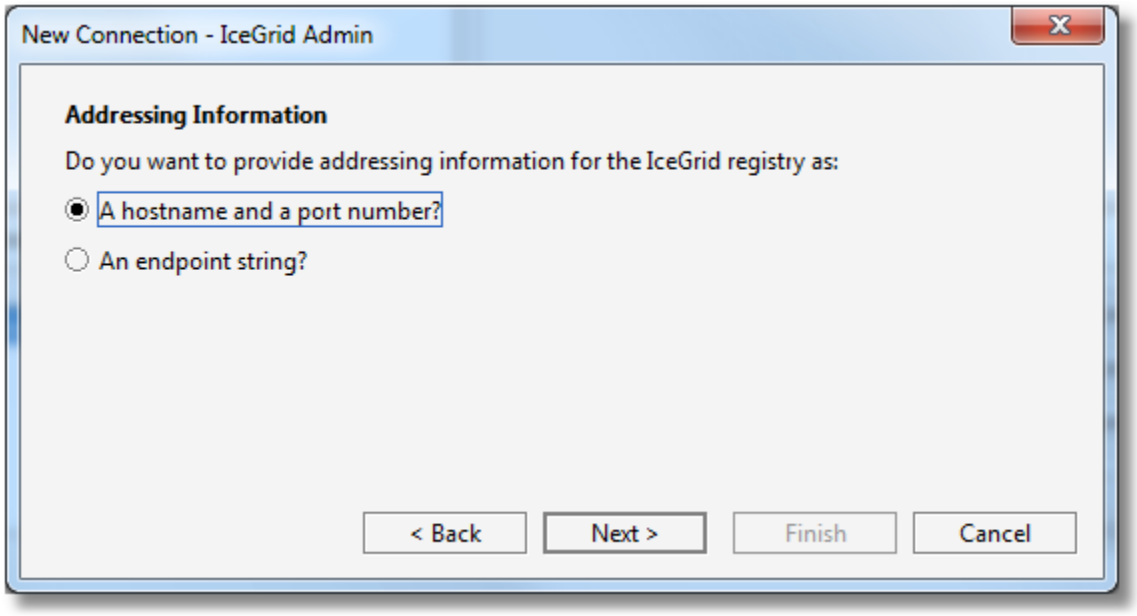

Step 5: Enter `localhost` for the hostname, leave the port number blank and keep TCP as the protocol:

Copyright 2016, ZeroC, Inc.

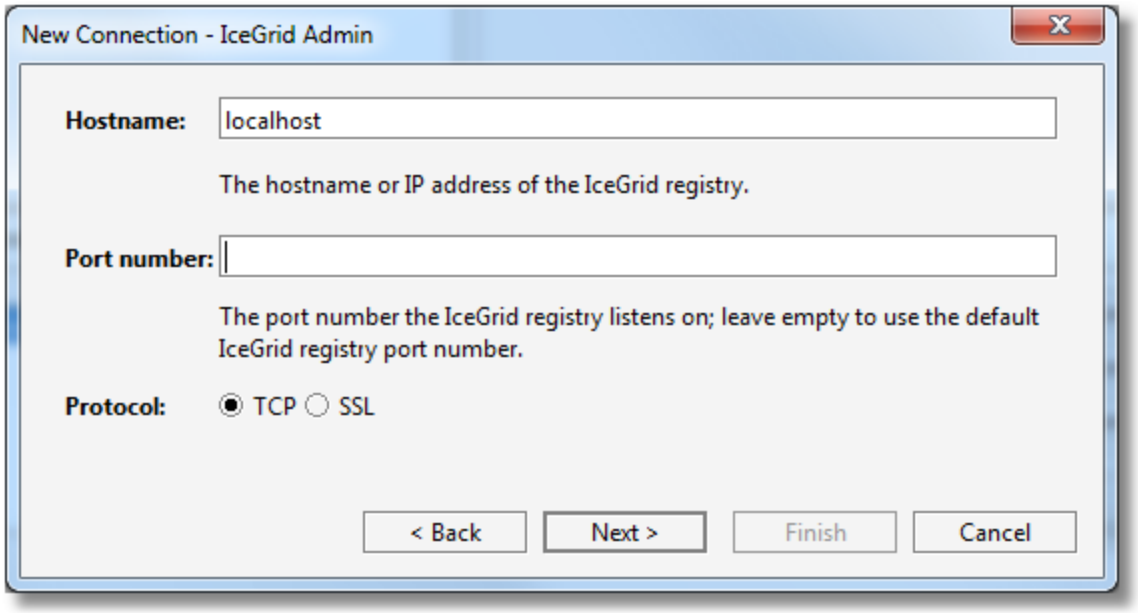

Step 6: Enter a username and password for this connection:

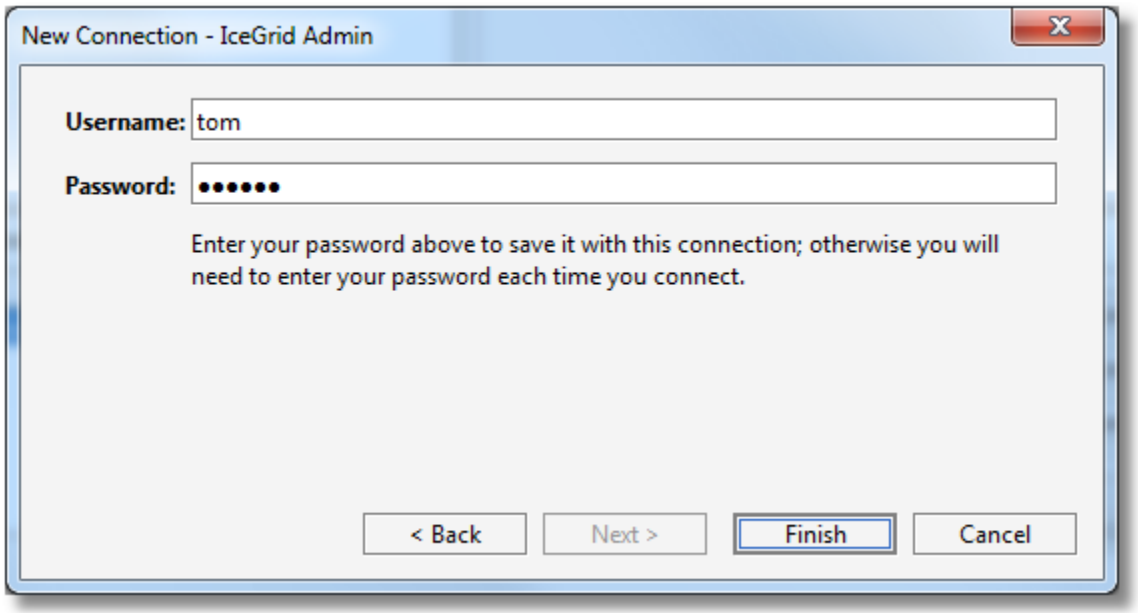

Step 7: Click `Finish` to save the connection; IceGrid Admin then attempts to connect to the IceGrid registry.

SSL Connection to local IceGrid using X.509 Credentials

Now let's create a direct SSL connection to an IceGrid registry and authenticate using our X.509 key (also used for SSL authentication). The target IceGrid registry must be configured to accept SSL connections and authentication using SSL credentials; see `IceGrid.Registry.AdminSSLPermissionsVerifier`.

Steps 1 to 4 are identical to the simple TCP connection described above.

Step 5: Enter the hostname of the IceGrid registry, leave the port number empty to use the default IceGrid port number (4061 for TCP and 4062 for SSL), and select SSL for the protocol:

Copyright 2016, ZeroC, Inc.

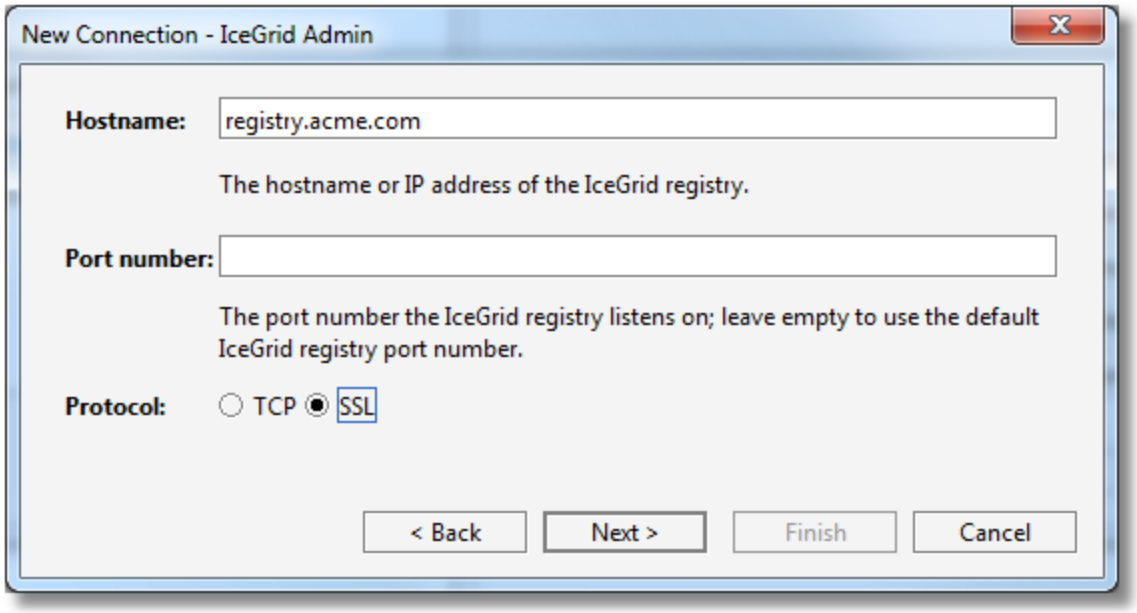

Step 6: Select Yes to provide an X.509 certificate for SSL authentication:

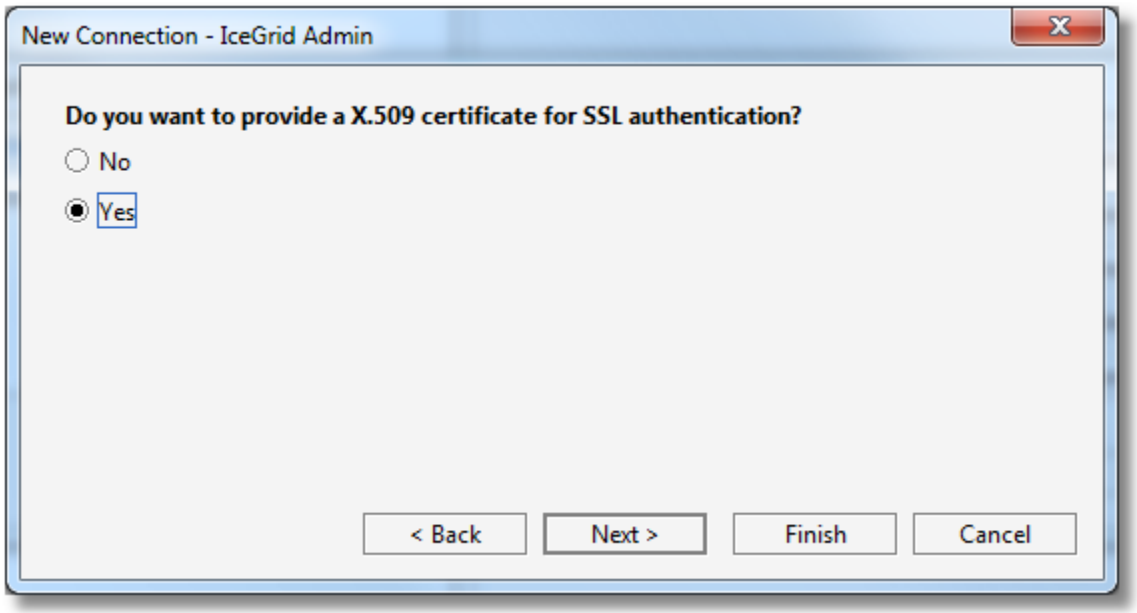

Step 7: Select the X.509 key used for SSL authentication from the Alias drop-down list; this list corresponds to the My Certificates set in the Certificate Manager described in the next section. Click on the `Import...` button to open the Certificate Manager dialog.

Copyright 2016, ZeroC, Inc.

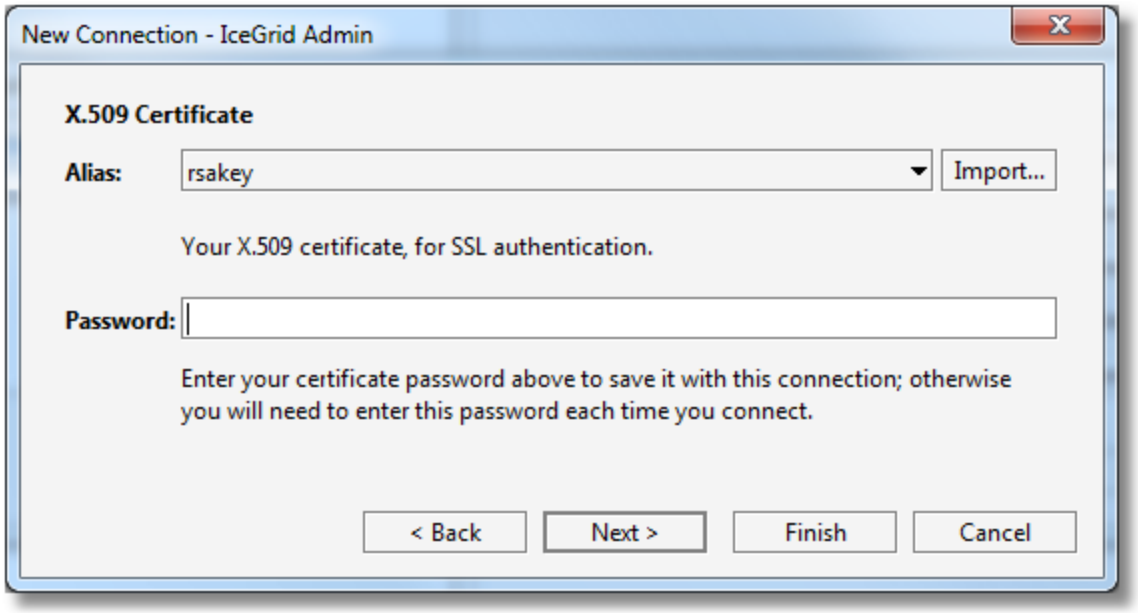

Step 8: Choose to use this X.509 certificate (carried through the SSL connection) to authenticate ourselves with the IceGrid registry:

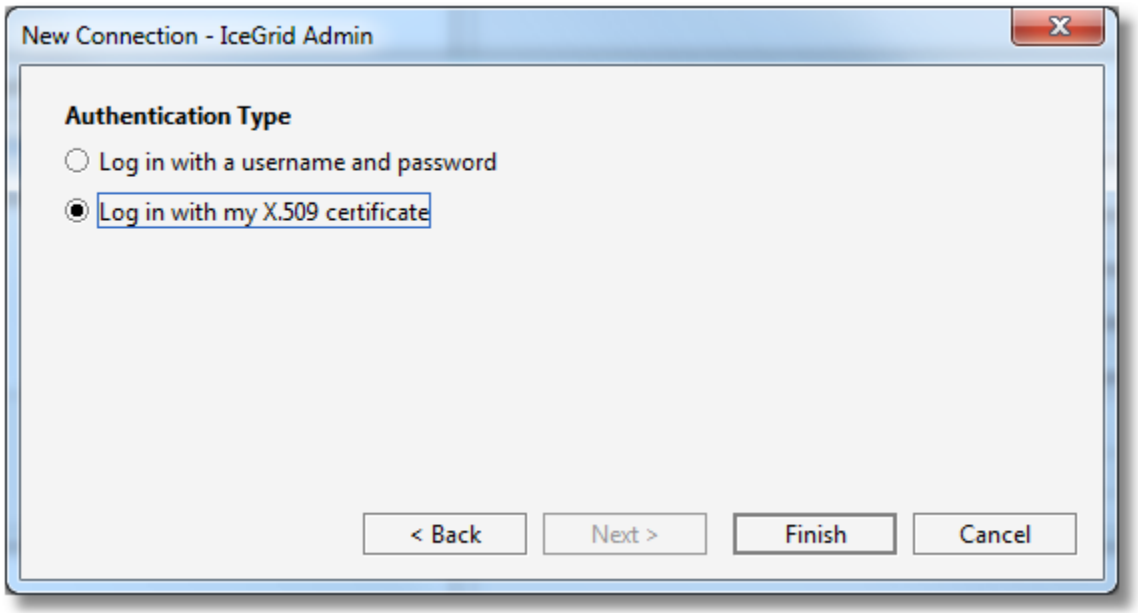

Step 9: Click `Finish` to save the connection and connect to the IceGrid registry. Unless we entered (and therefore saved) the X.509 key password with the connection, we are prompted for this password:

Copyright 2016, ZeroC, Inc.

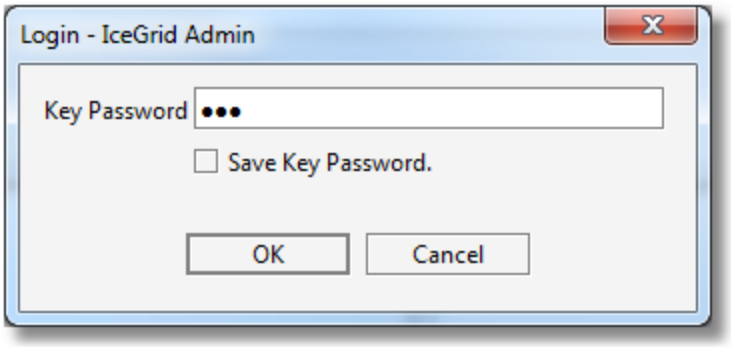

SSL Connection through Glacier2 router

Here we'll connect to an IceGrid registry "behind" a Glacier2 router. In this case, we need to connect to the Glacier2 router and authenticate ourselves with the router. We do not provide any information about the IceGrid registry itself: we will connect to the IceGrid registry identified by the target Glacier2 router configuration. See Glacier2 Integration with IceGrid.

Step 1: Select Routed Connection:

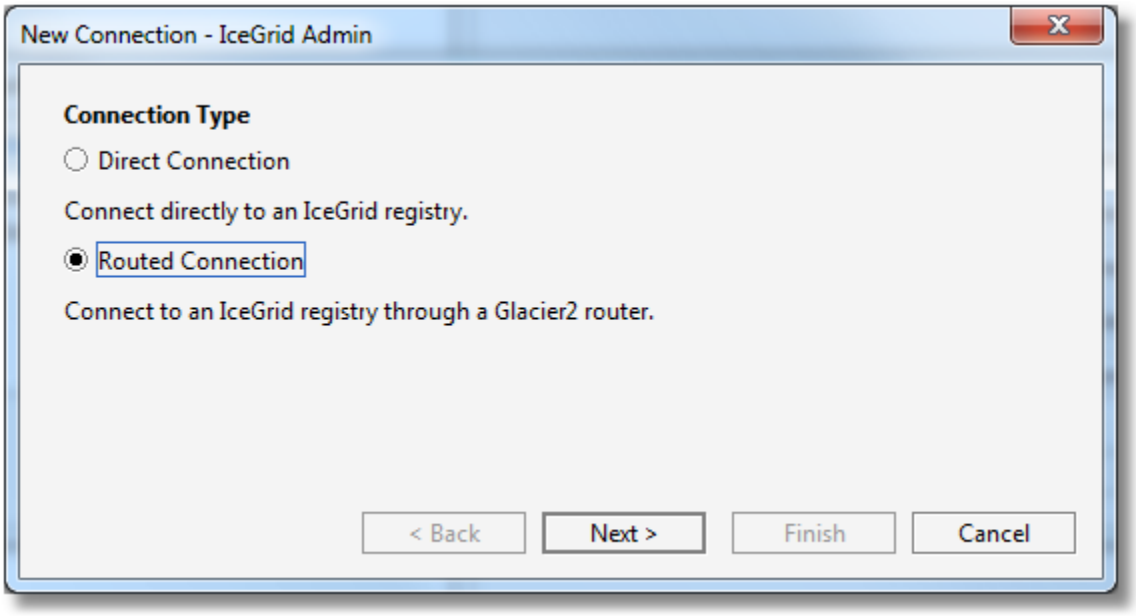

Step 2: Enter the addressing information for the target Glacier2 router as a hostname and port number:

Copyright 2016, ZeroC, Inc.

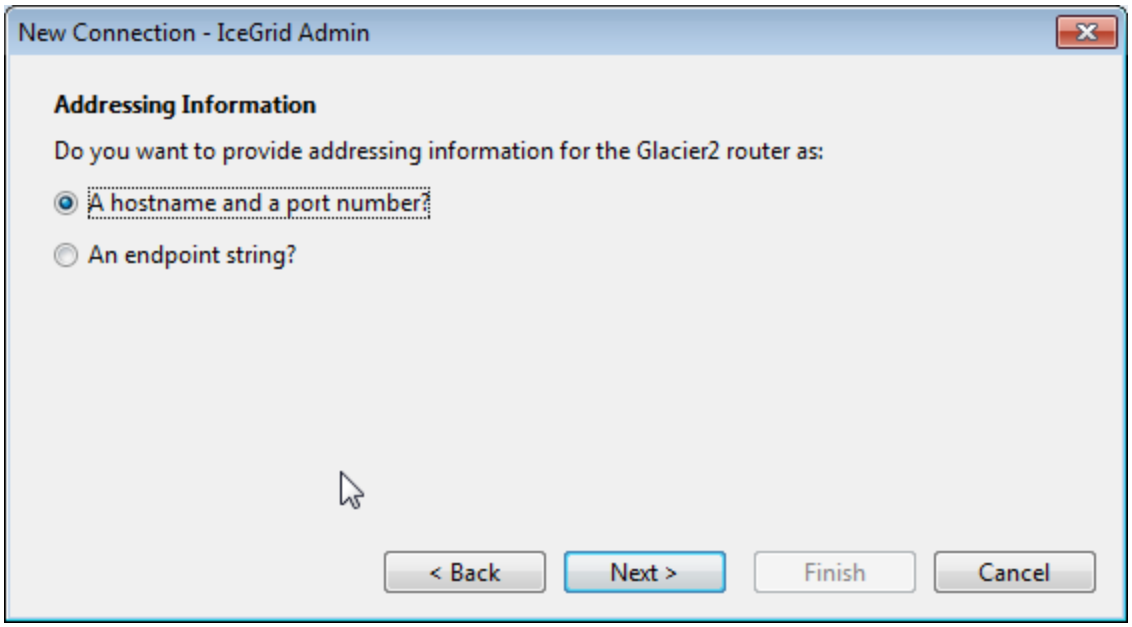

Step 3: Enter the hostname of the Glacier2 router, leave the port number empty to use the default Glacier2 port number (4063 for TCP and 4064 for SSL), and select SSL for the protocol:

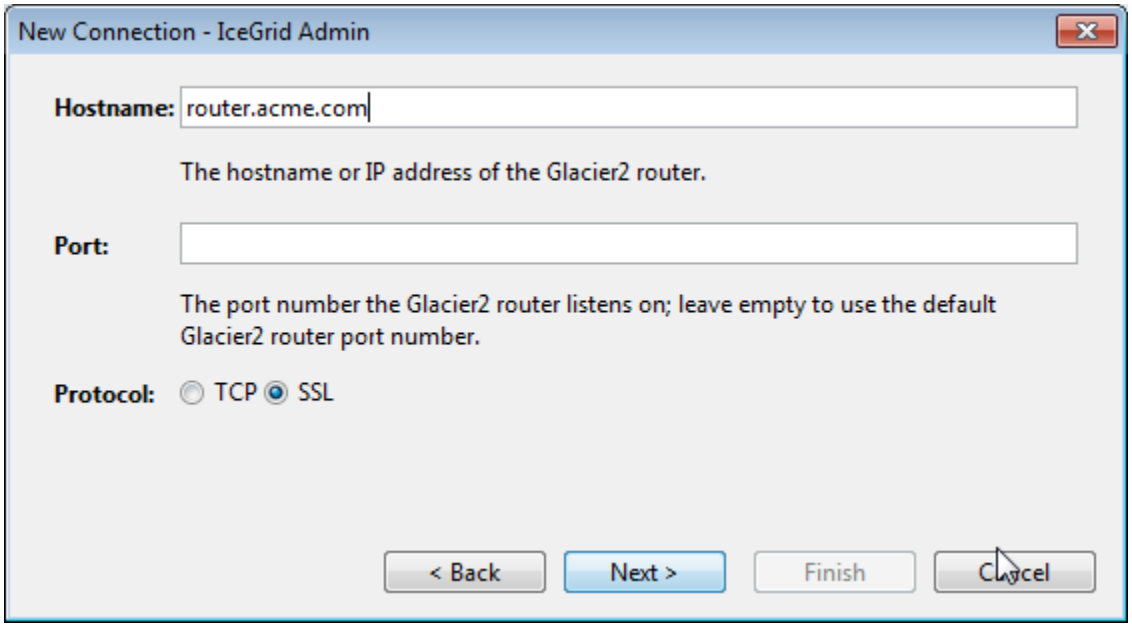

Step 4: Choose not to authenticate ourselves for SSL, so we do not provide a X.509 certificate:

Copyright 2016, ZeroC, Inc.

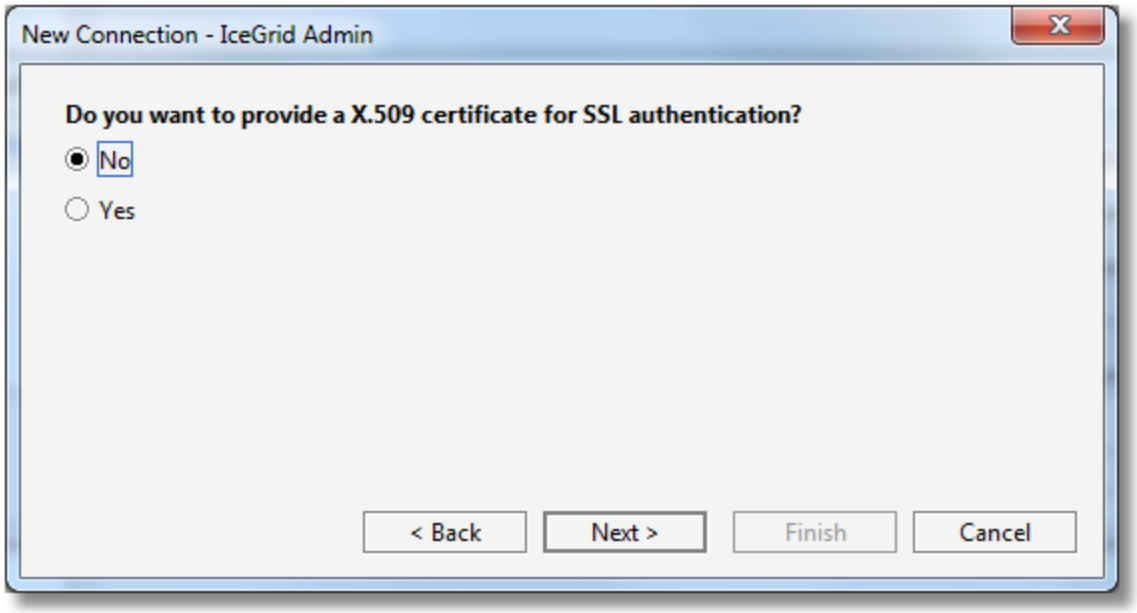

Step 5: Provide a username and password for the connection to the Glacier2 router:

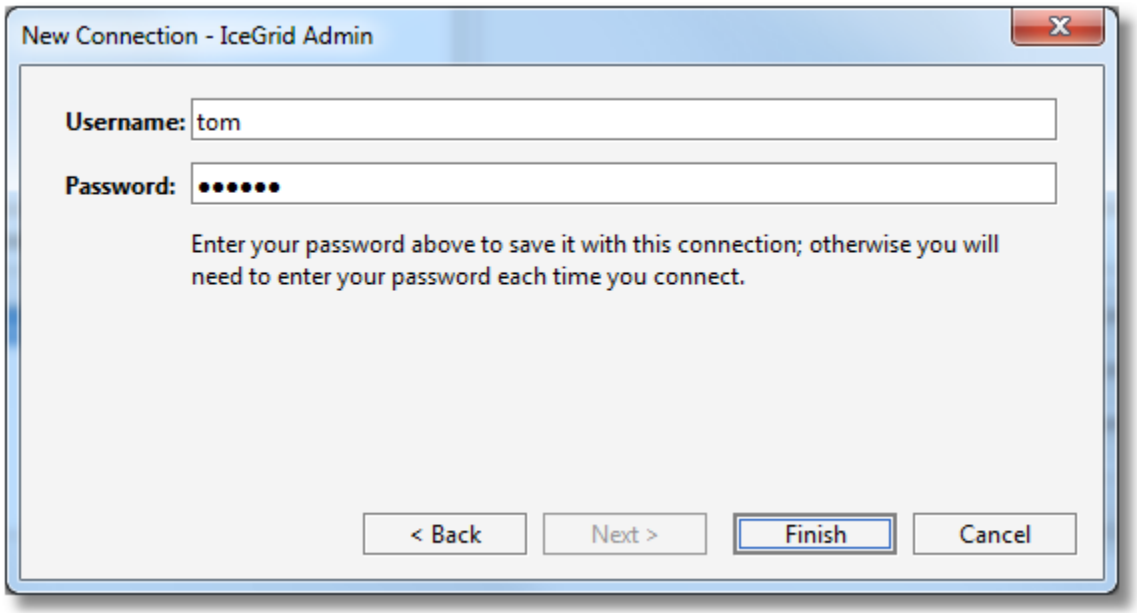

Step 6: Click `Finish` to save the connection; IceGrid Admin then attempts to connect to the IceGrid registry through the Glacier2 router.

SSL Connections and Certificates

IceGrid Admin maintains a persistent store of X.509 certificates for SSL connections with IceGrid registries. You can import, view and remove these certificates with the Certificate Manager. Use `File > Certificate Manager...` or click on the `Import...` button in the Connection wizard to open the Certificate Manager:

Copyright 2016, ZeroC, Inc.

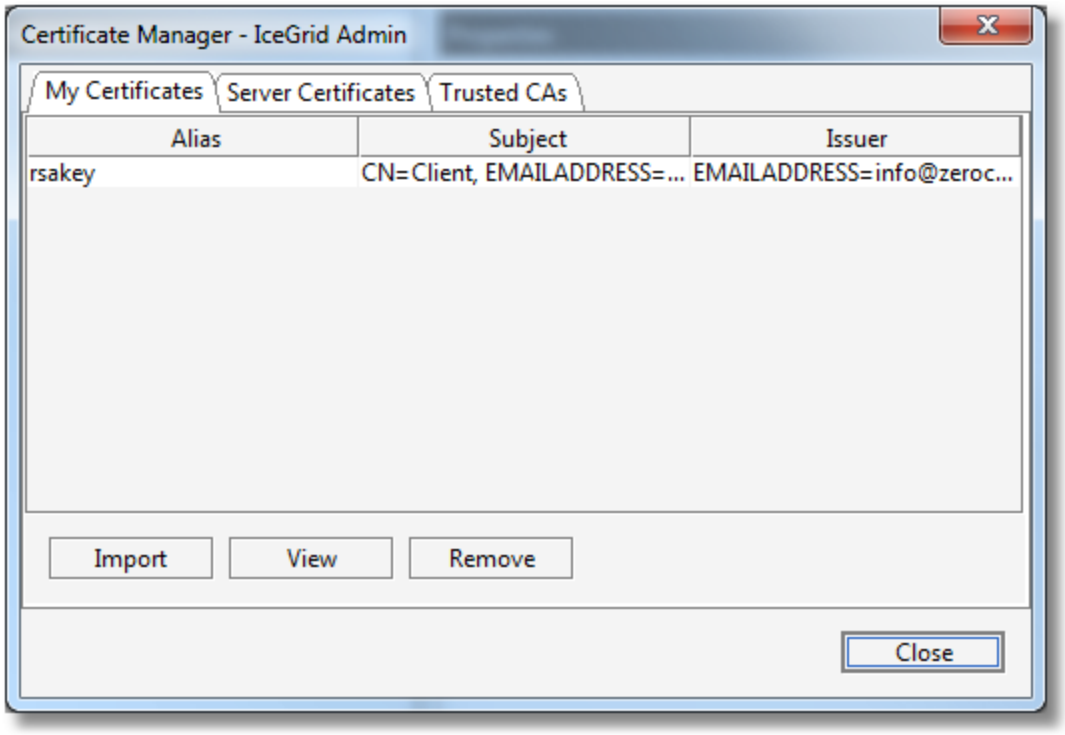

The Certificate Manager maintains three sets of certificates:

- **My Certificates**
  These certificates are used to authenticate IceGrid Admin when establishing an SSL connection with an IceGrid registry or Glacier2 router (this first authentication is at the SSL level); once the SSL connection is established, this certificate can also be used to authenticate with the target IceGrid registry or Glacier router (at the application level).

- **Server Certificates**
  The certificates of IceGrid registries and Glacier2 routers that IceGrid Admin trusts when establishing an SSL connection. All certificates signed by a Trusted CA (see below) are automatically trusted and usually do not need to be imported in this set.

- **Trusted CAs**
  The certificates of Certificate Authorities.

You do not need to import server certificates or CA certificates prior to establishing an SSL connection with an IceGrid registry or Glacier2 router. IceGrid Admin performs the following checks when establishing an SSL connection:

- If the X.509 certificate presented by the IceGrid registry or Glacier2 router matches a Server Certificate, proceed
- Otherwise, if this certificate is signed by a trusted CA, is valid (*now* is within the certificate's validity period) and its alternate name is a match for the SSL connection remote address, proceed
- Otherwise, display a Connection Security Warning dialog similar to the dialog below to let you decide whether or not to proceed with this certificate:

Copyright 2016, ZeroC, Inc.

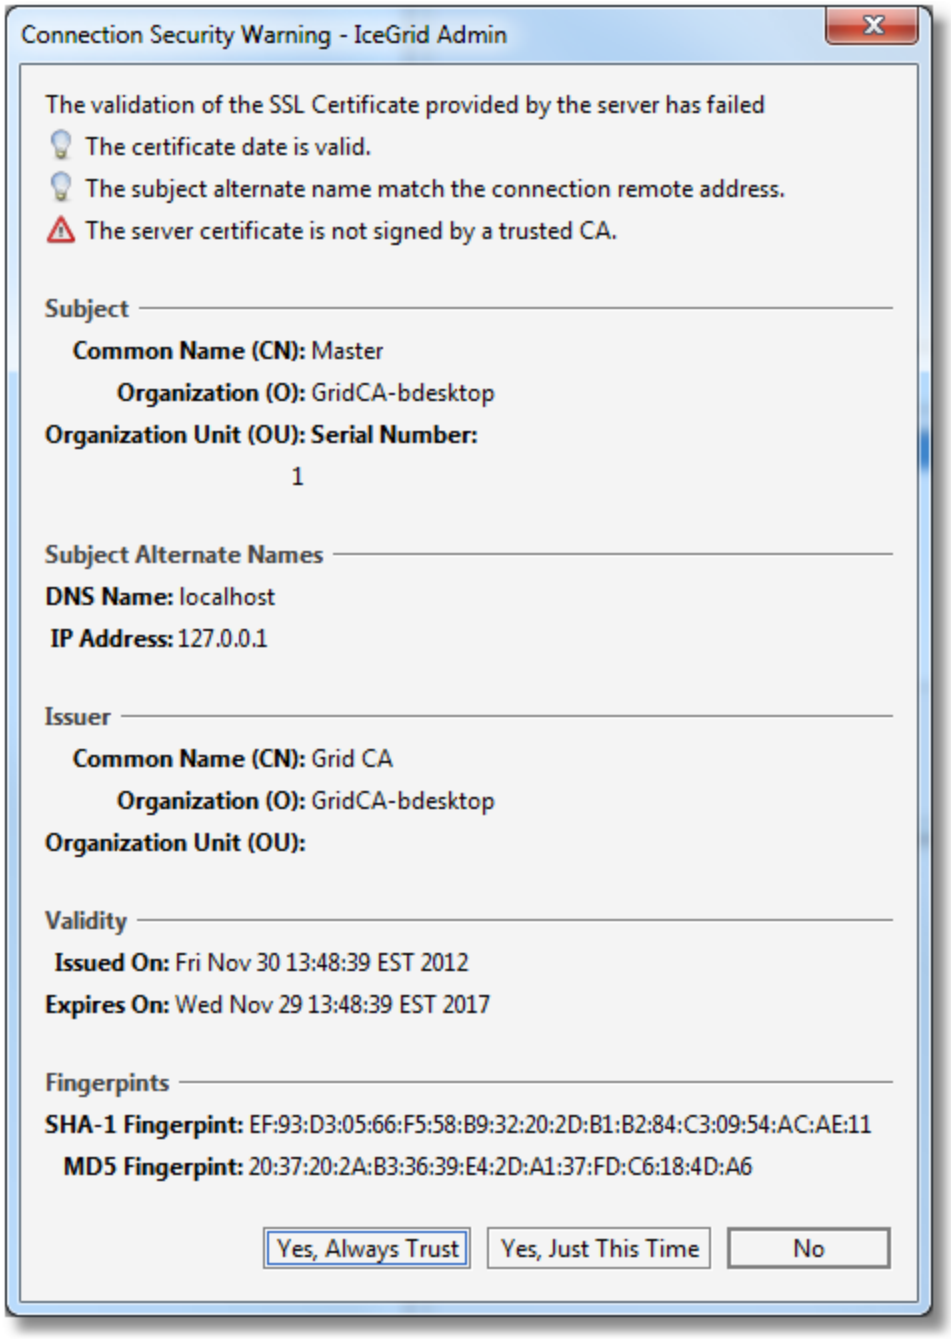

If you select `Yes, Always Trust`, the certificate is added in the persistent Server Certificates set.

A "client" X.509 certificate (saved in My Certificates) is only necessary when the target IceGrid registry or Glacier2 router requires one. This depends on the setting of the `IceSSL.VerifyPeer` property in those servers: when `IceSSL.VerifyPeer` is 2, IceGrid Admin must provide a valid certificate. If you forget to provide a certificate, or provide an invalid certificate, the connection establishment will fail with an IceSSL error comparable to the following:

Copyright 2016, ZeroC, Inc.

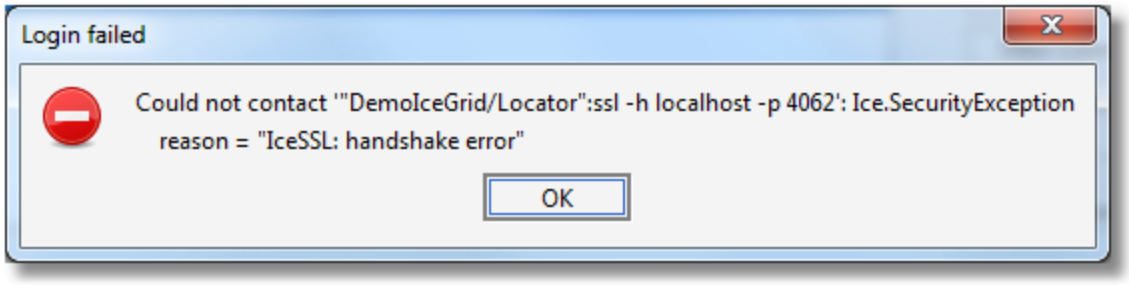

Editing a Saved Connection

In the `Saved Connections` dialog, click on `Edit Connection` to edit a connection. This opens the Connection wizard for your saved connection. With this wizard, IceGrid Admin does not attempt to connect to the target IceGrid registry when you click on the `Finish` button.

Closing a Connection

Use `File > Logout` or press the button to disconnect from an IceGrid registry. This clears all information in the Live Deployment pane.

Copyright 2016, ZeroC, Inc.

Run Time Components

IceGrid Admin shows the following run time components:

- Registry ( )
  Represents the IceGrid registry process with which IceGrid Admin communicates
- Slave Registry ( )
  Represents a slave registry process, when you have several replicated IceGrid registries
- Node ( )
  Represents an IceGrid node process
- Regular Server ( )
  Represents a regular server process, started and monitored by an IceGrid node
- IceBox Server ( )
  Represents an IceBox server process, started and monitored by an IceGrid node
- Adapter ( )
  Represents an Ice indirect object adapter, deployed within a server or a service
- Database Environment ( )
  Represents a Freeze/Berkeley DB Database Environment, deployed within a server or a service
- IceBox Service ( )
  Represents an IceBox service, deployed on an IceBox server
- Metrics View ()
  Represents a Metrics view associated with an Ice server or service.

All actions on these components (through menus, buttons and keyboard short-cuts) affect directly and immediately the running IceGrid deployment.

Copyright 2016, ZeroC, Inc.

Registry Run Time Component

The registry is the root node of the Run Time Components tree, and represents the IceGrid registry process administered by IceGrid Admin.

On this page:

- Actions
- Properties
- Children

Actions

A registry provides the following actions, from its contextual menu and from the `Tools > Registry` menu:

- **Add Well-Known Object**
  Create a new dynamic well-known object in the IceGrid registry.
- **Retrieve Ice log**
  Retrieve the log messages sent to the IceGrid registry's logger into an Ice Log Dialog. The Ice Log Dialog attaches a remote logger t o the registry's logger.
- **Retrieve stdout**
  Retrieve the IceGrid registry's stdout into a Log File Dialog. This retrieval succeeds only when the registry's stdout output has been redirected to a file using the `Ice.StdOut` property.
- **Retrieve stderr**
  Retrieve the IceGrid registry's stderr into a Log File Dialog. This retrieval succeeds only when the registry's stderr output has been redirected to a file using the `Ice.StdErr` property.
- **Shutdown**
  Shutdown the registry process.

You cannot restart an IceGrid registry from IceGrid Admin.

Properties

The Registry Properties panel shows:

- **Hostname**
  The name of the host on which the IceGrid registry process is running.
- **Deployed Applications**

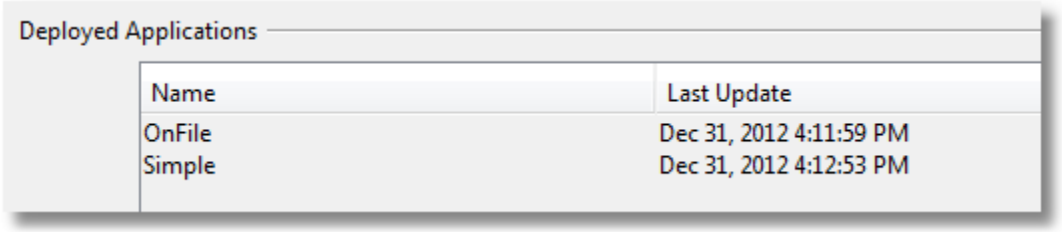

This table shows all the applications deployed on this IceGrid registry, along with the date and time of the last update of each application. A contextual menu allows you to:

  - open the corresponding application descriptor
  - patch the application, that is, instruct the IceGrid nodes to download the latest version of this application's files
  - show additional details on this application in this registry
  - remove (undeploy) the application from the registry

- **Dynamic Well-Known Objects**

Copyright 2016, ZeroC, Inc.

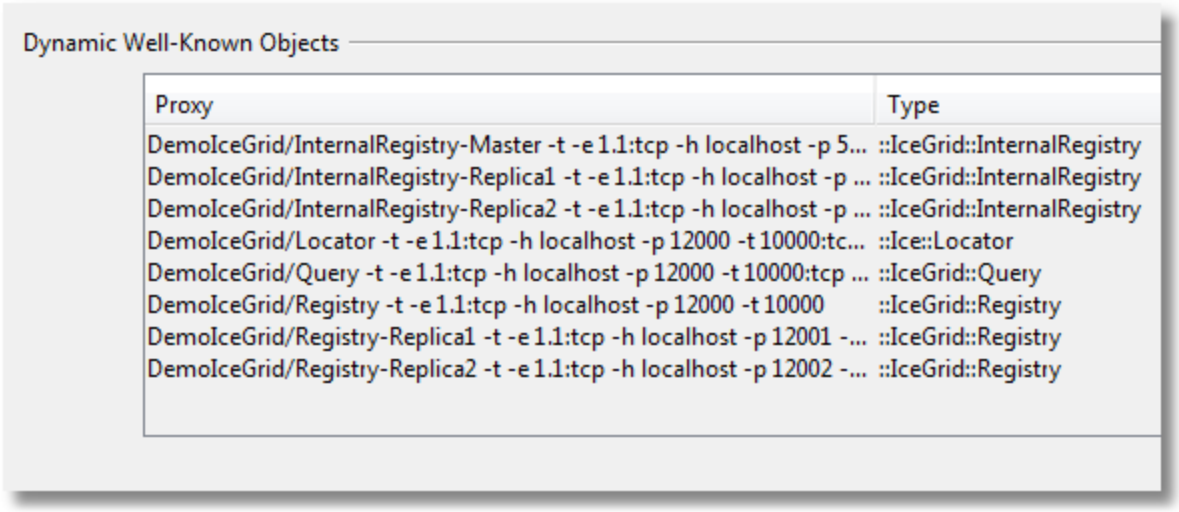

This table shows the well-known objects registry dynamically with the IceGrid registry: well-known objects defined using adapter and replica-group definitions are not included. A contextual menu allows you to add or remove entries from this table, and to show a given entry in its own dialog.

- **Dynamic Object Adapters**

  This table shows the object adapters registered dynamically with the registry. It is typically empty. A registry allows dynamically registered adapters only when its `IceGrid.Registry.DynamicRegistration` property is set to a value greater than 0. A contextual menu allows you to remove entries from this table.

  Note that application filtering does not affect this panel: all applications, well-known objects and dynamic object adapters are always displayed.

Children

A registry node can have two types of children:

- Slave Registry
- Node

Copyright 2016, ZeroC, Inc.

Slave Registry Run Time Component

An IceGrid deployment may use several IceGrid registry replicas, with one Master registry and a number of read-only Slave registries.

Actions

A slave registry provides the following actions, from its contextual menu and from the `Tools > Registry` menu

- **Retrieve Ice log**
  Retrieve the log messages sent to the IceGrid registry's logger into an Ice Log Dialog. The Ice Log Dialog attaches a remote logger t o the registry's logger.
- **Retrieve stdout**
  Retrieve the IceGrid registry's stdout into a Log File Dialog. This retrieval succeeds only when the registry's stdout output has been redirected to a file using the `Ice.StdOut` property.
- **Retrieve stderr**
  Retrieve the IceGrid registry's stderr into a Log File Dialog. This retrieval succeeds only when the registry's stderr output has been redirected to a file using the `Ice.StdErr` property.
- **Shutdown**
  Shutdown the registry process.

You cannot restart an IceGrid registry from IceGrid Admin.

Properties

The Registry Properties panel shows:

- **Hostname**
  The name of the host on which the IceGrid registry process is running.

Copyright 2016, ZeroC, Inc.

Node Run Time Component

A node represents an IceGrid node process registered with the IceGrid registry.

On this page:

- States
- Actions
- Properties
- Children

States

A node can be either up ( ) or down ( ). A "down" node is shown only when it is described by an application deployed on this IceGrid registry.

Actions

A node provides the following actions, from its contextual menu and from the `Tools > Node` menu:

- **Retrieve Ice log**
  Retrieve the log messages sent to the IceGrid node's logger into an Ice Log Dialog. The Ice Log Dialog attaches a remote logger to the node's logger.
- **Retrieve stdout**
  Retrieve the IceGrid node's stdout into a Log File Dialog. This retrieval succeeds only when the node's stdout output has been redirected to a file using the `Ice.StdOut` property.
- **Retrieve stderr**
  Retrieve the IceGrid node's stderr into a Log File Dialog. This retrieval succeeds only when the node's stderr output has been redirected to a file using the `Ice.StdErr` property
- **Shutdown**
  Shutdown the IceGrid node process.

You cannot restart an IceGrid node from IceGrid Admin.

Properties

Copyright 2016, ZeroC, Inc.

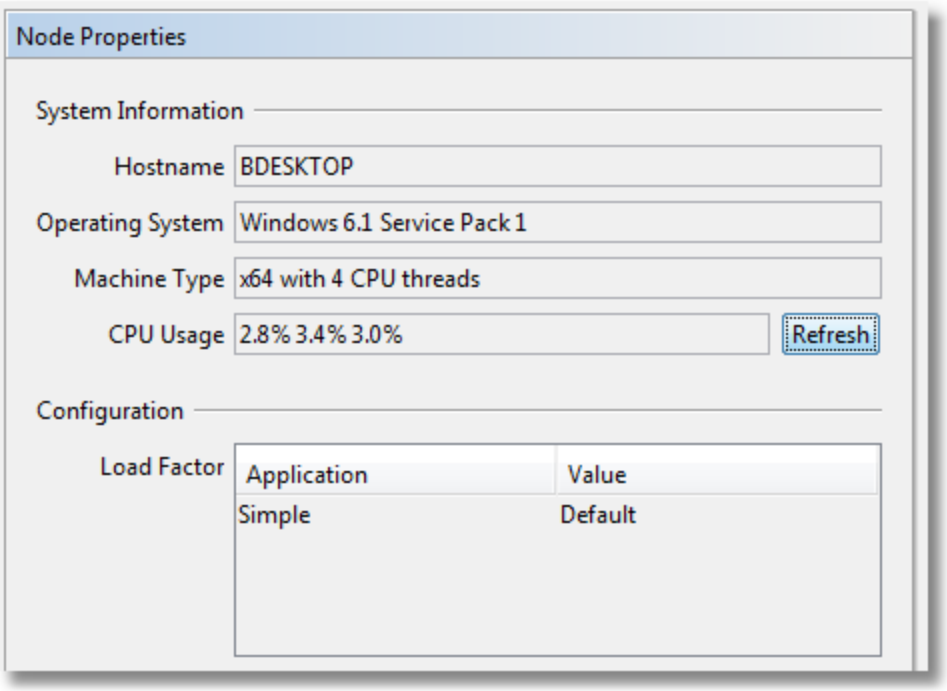

The Node Properties panel shows:

- **Hostname**
  The name of the host on which the IceGrid node process is running.
- **Operating System**
  The operating system name and version of the host on which the IceGrid node process is running.
- **Machine Type**
  The type of CPUs and the number of CPU threads of the host on which the IceGrid node process is running. For example, a computer with a dual-core CPU and 2 "hyper threads" per core will show 4 CPU threads.
- **CPU Usage (Windows) or Load Average (Linux/Unix)**
  On Windows, shows the percentage of CPU utilization in the past 1, 5, and 15 minutes. On Linux/Unix, shows the load-average in the past 1, 5 and 15 minutes. These values are retrieved when the Node Properties panel is displayed. Click on the Refresh button to retrieve the latest values.
- **Load Factor**
  Shows the load factor defined by each application using this node.

Children

A registry node can only have two types of children:

- Regular Server
- IceBox Server

Copyright 2016, ZeroC, Inc.

Server Run Time Component

A server represents an Ice server process. It can be either regular server (with typically a single Ice communicator) or an IceBox server hosting a number of IceBox services.

On this page:

- States
- Actions
- Properties
- Children

States

A server is always in one of the following states (the first icon is for regular servers, the second for IceBox servers):

- Unknown ( ): this state is shown when the parent IceGrid node is down.
- Inactive ( ): the server is not running.
- Activating ( ): the server is starting up. The IceGrid registry is waiting for the server to register all its object adapters with server lifetime.
- Active ( ): the server is running, and has registered all its object adapters with server lifetime with the IceGrid registry.
- Deactivating ( ): the server is shutting down. The IceGrid registry is waiting for the server process to exit.
- Destroyed ( ): the server being removed of the IceGrid registry. This is a very transient state.

A server can also be either enabled or disabled; when disabled, the icons above are grayed-out. A disabled server cannot be started until it is re-enabled.

Actions

A server provides the following actions, from its contextual menu, from the `Tools > Server` menu, and from buttons on the Server Properties panel:

- **Start**
  Instruct the IceGrid node to start the server.

- **Stop**
  Instruct the IceGrid node to shutdown the server.

- **Enable**
  Mark the server as "enabled".

- **Disable**
  Mark the server as "disabled". A disabled server cannot be started; however an already running server can be marked "disabled".

- **Patch Distribution**
  Patch the server's distribution, that is, instruct the IceGrid node to download the latest files from the server's distribution.

- **Write Message**
  Open a dialog that allows you to write a message to the server's stdout or stderr.

- **Retrieve Ice log**
  Retrieve the log messages sent to the server's logger into an Ice Log Dialog. The Ice Log Dialog attaches a remote logger to the server's logger.

- **Retrieve stdout**
  Retrieve the stdout log file of this server into a Log File Dialog. This retrieval succeeds only when the server's stdout output has been redirected to a file using the `Ice.StdOut` property. This is usually achieved by setting the IceGrid.Node.Output property in the IceGrid node configuration file.

- **Retrieve stderr**
  Retrieve the stderr log file of this server into a Log File Dialog. This retrieval succeeds only when the server's stderr output has been redirected to a file using the `Ice.StdErr` property. This is usually achieved by setting the IceGrid.Node.Output property in the IceGrid node configuration file.

Copyright 2016, ZeroC, Inc.

- **Retrieve log file**
  Retrieve a log file of this server into a Log File Dialog.

- **Send Signal**
  Send a signal to a server, for example SIGQUIT. Available only for non-Windows servers.

Properties

The Server Properties panel shows first the Runtime Status of the server, i.e. "live" values retrieved from the server:

- **State**
  The state of the server (Active, Deactivating, Inactive etc., see above)

- **Enabled**
  A checkbox that is checked when the server is enabled.

- **Process Id**
  The process ID of the server.

- **Build Id**
  The build Id of this server: this corresponds to the Ice property BuildId.

- **Properties**
  A table showing all the Ice properties currently set in this server. These properties are retrieved each time you select a new server in IceGrid Admin, and each time you click on the Refresh button next to the Build Id field.

The remaining Server Properties under Configuration come from the IceGrid descriptors associated with this server:

- **Application**

  The name of the application containing this server's definition. The button on the right shows the server definition in an Application tab.

- **Description**
  A free-text description of this server.

- **Properties**
  A table showing all the Ice properties of this server. These properties may come from template definitions, property sets, server-instance properties etc. They are all combined in this table.

- **Path to Executable**
  The path to the server's executable, used by the IceGrid node to start the server.

- **Ice Version**
  The Ice version of this server.

- **Working Directory**
  The path to the server's working directory used by the IceGrid node when starting the server.

- **Command Arguments**
  The command-line arguments given to the server when started by IceGrid.

- **Run as**
  On Linux/Unix, a server may be started as a specific user when IceGrid node runs as root. Run as shows this username. When blank, the server runs as the IceGrid node user except if IceGrid node runs as root (on Linux/Unix); in this case, the server runs as nobody.

- **Environment Variables**
  A server started by IceGrid node gets these environment variables in addition to the environment variables inherited from the IceGrid node.

- **Activation Mode**

Copyright 2016, ZeroC, Inc.

The server's activation mode.

- **Activation Timeout**
  The server's activation timeout.

- **Deactivation Timeout**
  The server's deactivation timeout.

- **Allocatable**
  This checkbox Shows whether this server is allocatable or not.

- **Depends on the application distribution**
  If this checkbox is checked, this server depends on the enclosing application's distribution: each time this distribution is patched, the server is automatically shut down before the patch.

- **IcePatch2 Proxy**
  A stringified or well-known proxy for the IcePatch2 server than contains this server's distribution. When blank, this server does not have its own distribution.

- **Directories**
  List of directories included in the server distribution. When blank, the entire IcePatch2 server repository is used as the server distribution.

Children

A regular server node can have two types of children:

- Adapter
- Database Environment

An IceBox server node can have only one type of children:

- Service

Copyright 2016, ZeroC, Inc.

Adapter Run Time Component

An adapter represents an Ice object adapter described in the IceGrid registry that resides in a server or an IceBox service. Note that direct object adapters are not displayed since IceGrid knows nothing about them.

On this page:

- States
- Properties

States

An adapter can be either active ( ) or inactive ( ).

Properties

The Adapter Properties panel shows first the Runtime Status of the object adapter, i.e. "live" values retrieved from the enclosing server or service:

- **Status**
  The status of this object adapter: Active or Inactive.

- **Published Endpoints**
  The endpoints that this object adapter has registered with the IceGrid registry. Blank when the adapter's status is Inactive.

The remaining Server Properties under Configuration come from the IceGrid descriptor that defined this object adapter:

- **Description**
  A free-text description of this adapter.

- **Adapter ID**
  This object adapter's ID. Adapter IDs are unique within an IceGrid deployment.

- **Replica Group**
  When this adapter belongs to a replica group, shows the replica group ID. Blank otherwise.

- **Priority**
  The adapter's priority. Used only when the adapter belongs to a replica group with the "Ordered" load balancing policy.

- **Endpoints**
  The Endpoints configuration property of this object adapter.

- **Published Endpoints**
  The PublishedEndpoints configuration property of this object adapter.

- **Register Process**
  When the Register Process checkbox is checked, this object adapter will register an `Ice::Process` object with IceGrid upon activation. This setting is meaningful only for servers with an Ice version less than 3.3.

- **Server Lifetime**
  When the Server Lifetime checkbox is checked, IceGrid considers that this object adapter is activated when the server starts up and deactivated when the server shuts down. This allows IceGrid to detect when a server is fully activated (all its object adapters with a server-lifetime have registered) and when a server is shutting down (at least one object adapter with server-lifetime has unregistered).

- **Well-known Objects**
  The well-known objects defined by this object adapter.

- **Allocatable Objects**
  The allocatable objects defined by this object adapter.

Copyright 2016, ZeroC, Inc.

Database Environment Run Time Component

A Database Environment ( ) represents a Berkeley DB database environment, used by the Freeze persistence service.

Properties

The Database Environment Properties panel shows:

- **Description**
  A free-text description of this Berkeley DB database environment.
- **DB Home**
  The path to the home directory of this database environment, which corresponds to the Ice property `Freeze.DbEnv.`*`env-name`*`.DbHome`. It is often created by the IceGrid node in the server's private directory.
- **Properties**
  Berkeley DB configuration properties for this database environment. IceGrid generates a `DB_CONFIG` file in the DB Home directory with these properties.

Copyright 2016, ZeroC, Inc.

Service Run Time Component

A service represents an IceBox service loaded (or potentially loaded) within an IceBox server.

On this page:

- States
- Actions
- Properties
- Children

States

A service can be either started ( ) or stopped ( ) within an IceBox server.

Actions

An IceBox service provides the following actions, from its contextual menu, from the `Tools > Service` menu, and from buttons on the Service Properties panel:

- **Start**
  Instruct the IceBox server to start the service.

- **Stop**
  Instruct the IceBox server to stop this service.

- **Retrieve Ice log**
  Retrieve the log messages sent to the service's logger into an Ice Log Dialog. The Ice Log Dialog attaches a remote logger to the service's logger.

- **Retrieve log file**
  Retrieve the log file of this service into a Log File Dialog.

Properties

The Service Properties panel shows first the Runtime Status of the service, i.e. "live" values retrieved directly from the service:

- **State**
  A checkbox that is checked when the service is started.

- **Build Id**
  The build Id of this service: this corresponds to the Ice property `BuildId`.

- **Properties**
  A table showing all the Ice properties currently set in this service. These properties are retrieved each time you select a new service in IceGrid Admin, and each time you click on the Refresh button next to the Build Id field.

The remaining Server Properties come from the IceGrid descriptors associated with this service:

- **Description**
  A free-text description of this service.

- **Properties**
  A table showing all the Ice properties of this service. These properties may come from template definitions, property sets, service-instance properties etc. They are all combined in this table.

- **Entry Point**
  The entry point for this service. This corresponds to the value of the `IceBox.Service.`*`name`* property.

Children

An IceBox service can have two types of children:

- Adapter
- Database Environment

Copyright 2016, ZeroC, Inc.

Copyright 2016, ZeroC, Inc.

Metrics View Run Time Component

A Metrics View displays the Metrics maps associated with a Server or Service.

On this page:

- States
- Actions
- Metrics Report

States

A metrics view can be either enabled ( ) or disabled( ). Once enabled, a Metrics View may degrade the performance of the instrumented server or service.

Actions

An Metrics View provides the following actions, from its contextual menu and from the `Tools > Metrics View` menu:

- **Enable**
  Enable this Metrics View.

- **Disable**
  Disable this Metrics View.

Metrics Report

The Metrics Report panel shows the maps included in the Metrics View. The columns of the maps include the metrics themselves (for example, the total number of operations dispatched by the server since the Metrics view was enabled) and computed values (for example, the average lifetime of an operation dispatch, since the Metrics view was enabled).

Tool tips on each column describe the metrics or computed value displayed by the column.

IceGrid Admin retrieves the latest metrics every 5 seconds and automatically refreshes the current Metrics Report.

Note that IceGrid Admin displays only non-empty metrics maps; for example, if a server does not make any remote invocation, its Invocation map is empty and will not be displayed.

Copyright 2016, ZeroC, Inc.

Application Component

An application is primarily a way to group definitions and does not actually exist in a live deployment.

Nevertheless, IceGrid Admin offers a number of application-related actions in the live deployment view through the Filter live deployment button and the `Tools > Application` menu:

- **Patch Distribution**
  Instruct the IceGrid nodes to download the latest version of the files associated with the selected application.

- **Show details**
  Show various information about the selected application.

- **Filter live deployment**
  Shows only the selected application in the live deployment view; servers and other components associated with other applications are "filtered out".

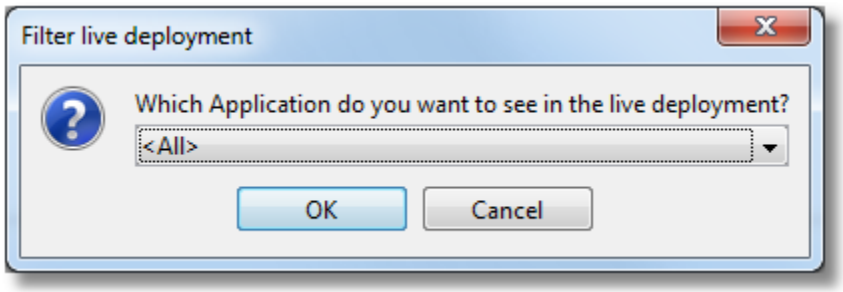

  When filtering is enabled, the name of the selected application appears to the right of the registry name. In the picture below, the live deployment view shows only the application named 'Simple'.

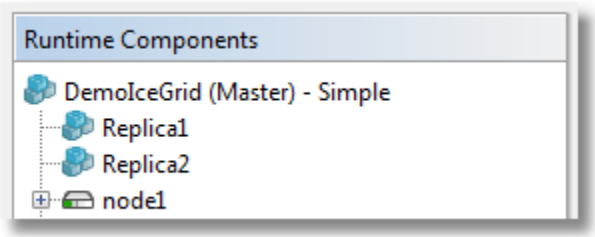

- **Remove from registry**
  Remove the selected application from the registry.

Copyright 2016, ZeroC, Inc.

Ice Log Dialog

The Ice Log dialog shows the Ice log messages of a server, service, node or registry retrieved through IceGrid. The monitored Ice application sends these log messages to a remote logger implemented by the Ice Log dialog.

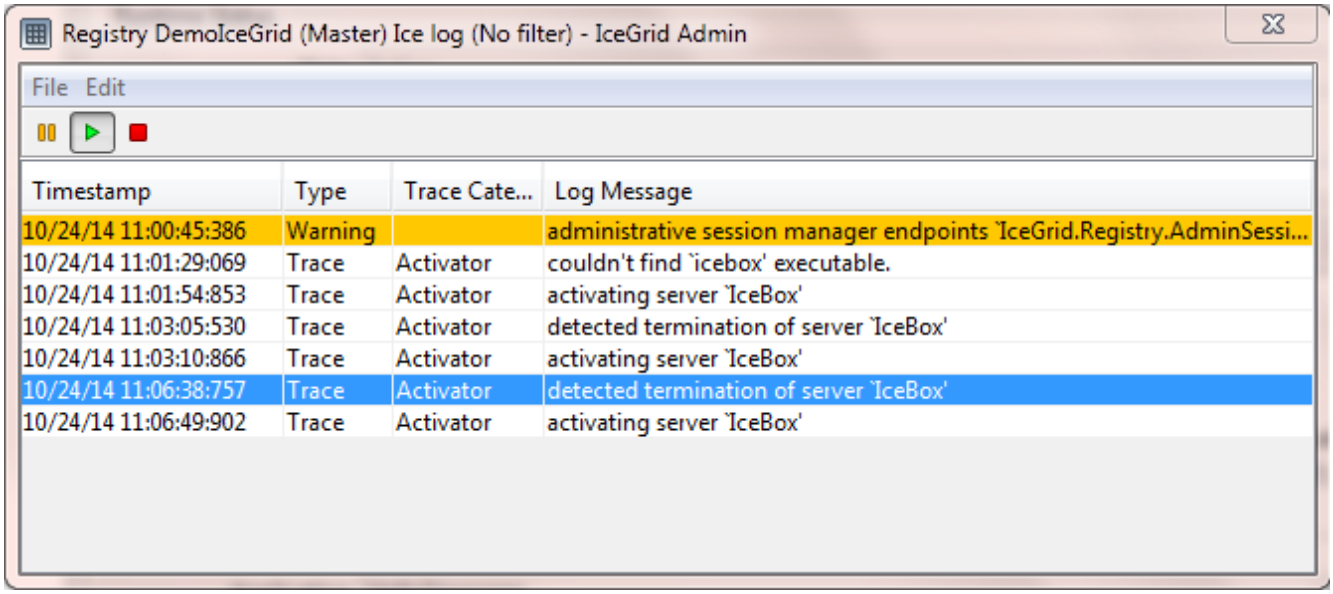

On this page:

- States
- Preferences
- Filter

States

An Ice log dialog is always in one of the following states:

- **Running**
  The dialog is attached to the monitored application's logger: it displays new log messages as soon as it receives them.

- **Paused**
  The dialog is attached to the monitored application's logger: it queues new log messages, but does not display them.

- **Stopped**
  The dialog is no longer attached to the monitored application's logger.

When an Ice log dialog is running or paused, and the target server, node or registry is stopped, the Ice log dialog automatically transitions to the Stopped state. However, when an Ice log dialog is attached to an IceBox service, the Ice log dialog's state does not change when only the service is stopped, because stopping a service does not shutdown or destroy the associated communicator.

Preferences

Use the `Edit > Preferences...` menu to open the Preferences dialog. These preferences apply to the current dialog and to any Ice Log dialog opened later on.

Copyright 2016, ZeroC, Inc.

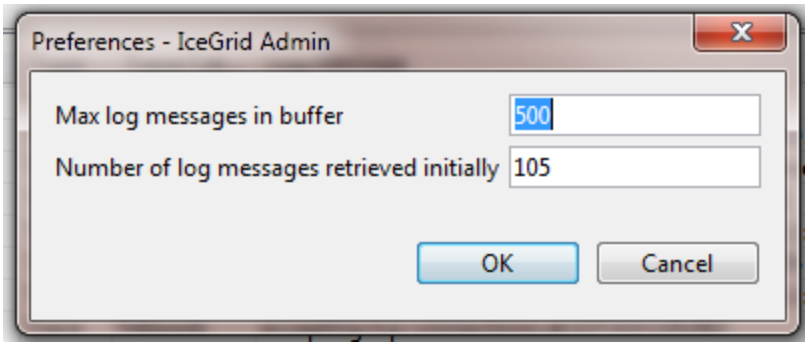

This dialog allows you to view and update the following settings:

- **Max log messages in buffer**
  The maximum number of lines displayed in the Ice log dialog.
- **Number of log messages retrieved initially**
  When a new dialog is opened, or when restarting a stopped dialog, the dialog retrieves and displays up to this number of lines.

Filter

Use the `Edit > Filter...` menu to open the Filter dialog. This filter applies only to the current Ice Log dialog.

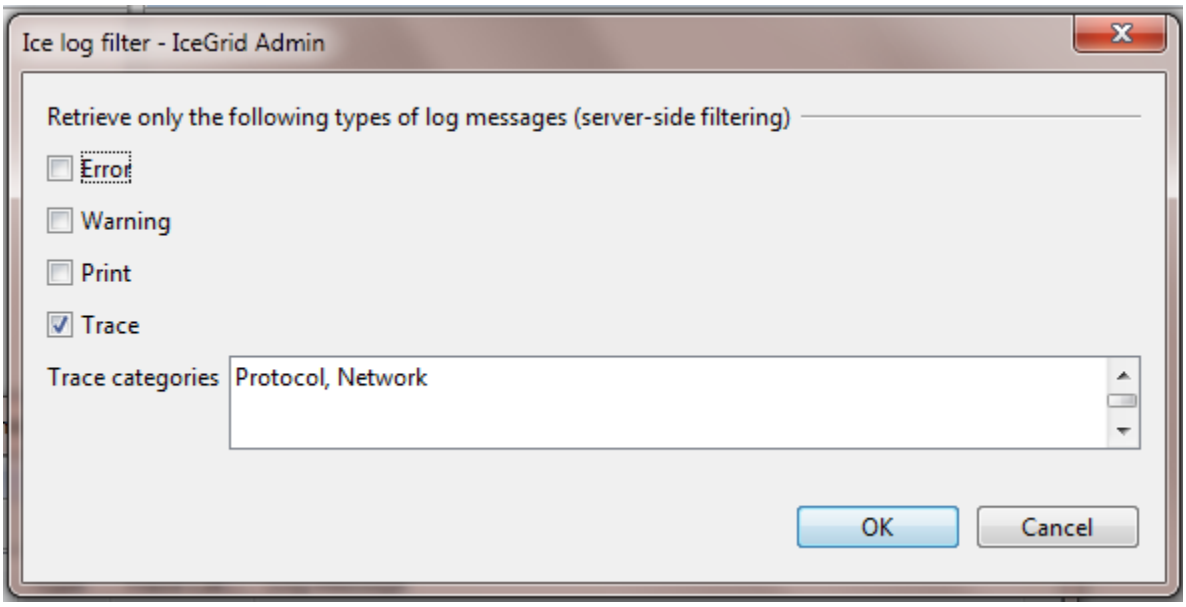

This dialog allows you to filter the log messages sent by the monitored Ice applications. All the filtering is done in this remote Ice application.

- **Error, Warning, Print and Trace Check Boxes**
  Check the type of log messages you want to receive.
- **Trace Categories**
  Enter the list of trace categories you want to receive, or leave empty to receive trace messages with any trace category.

Copyright 2016, ZeroC, Inc.

Log File Dialog

The Log File dialog shows a log file (text file) retrieved through IceGrid. Often, a process will be writing to this log file and the dialog will periodically check for new lines to display.

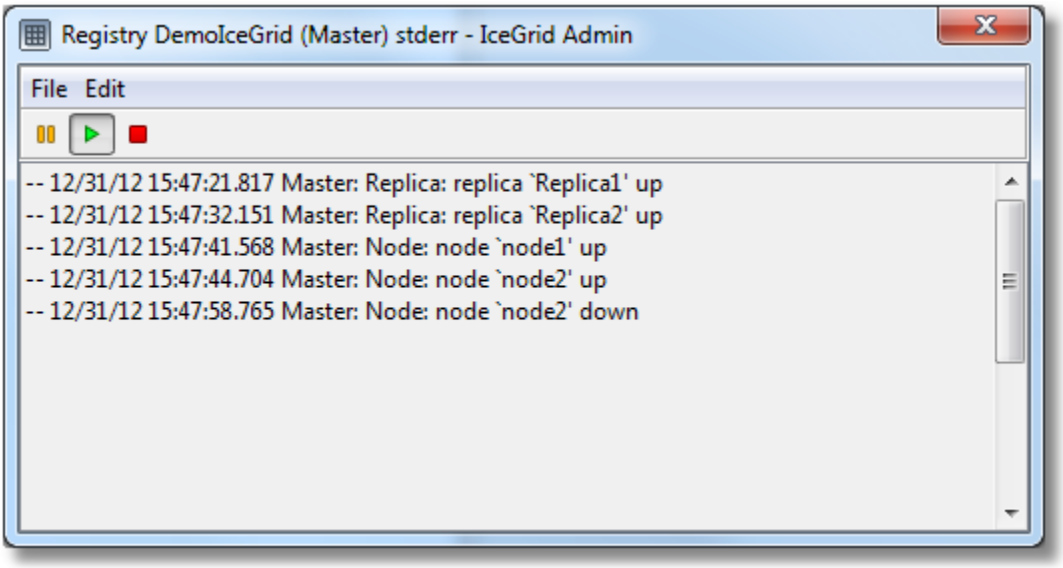

On this page:

- States
- Preferences

States

A log file dialog is always in one of the following states:

- **Running**
  The dialog is periodically checking for new lines.

- **Paused**
  The dialog is "connected" to a remote log file but does not currently retrieve new lines.

- **Stopped**
  The dialog is not retrieving new lines from the remote log file.

Preferences

Use the `Edit > Preferences...` menu to open the Preferences dialog. These preferences apply to the current dialog and to any Log File dialog opened later on.

Copyright 2016, ZeroC, Inc.

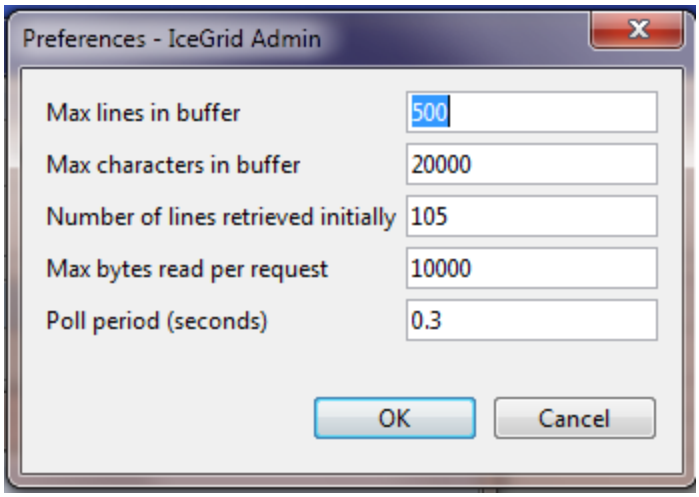

This dialog allows you to view and update the following settings:

- **Max lines in buffer**
  The maximum number of lines displayed in the log dialog.
- **Max characters in buffer**
  The maximum number of characters displayed in the log dialog.
- **Number of lines retrieved initially**
  When a new dialog is opened, or when restarting a stopped dialog, the dialog retrieves and displays up to this number of lines.
- **Max bytes read per request**
  The maximum number of bytes retrieve by each request. Pick a value that is low enough to make the dialog appear responsive and big enough to avoid many round-trips when lots of data are logged. IceGrid Admin requires a value between 100 and Ice.MessageSizeMax - 512.
- **Poll period**
  When in the running state, the dialog attempts to retrieve new lines from the log file (through IceGrid) every Poll period seconds.

Copyright 2016, ZeroC, Inc.

Metrics Graph

A Metrics Graph displays metrics retrieved from one or more Ice servers or IceBox services.

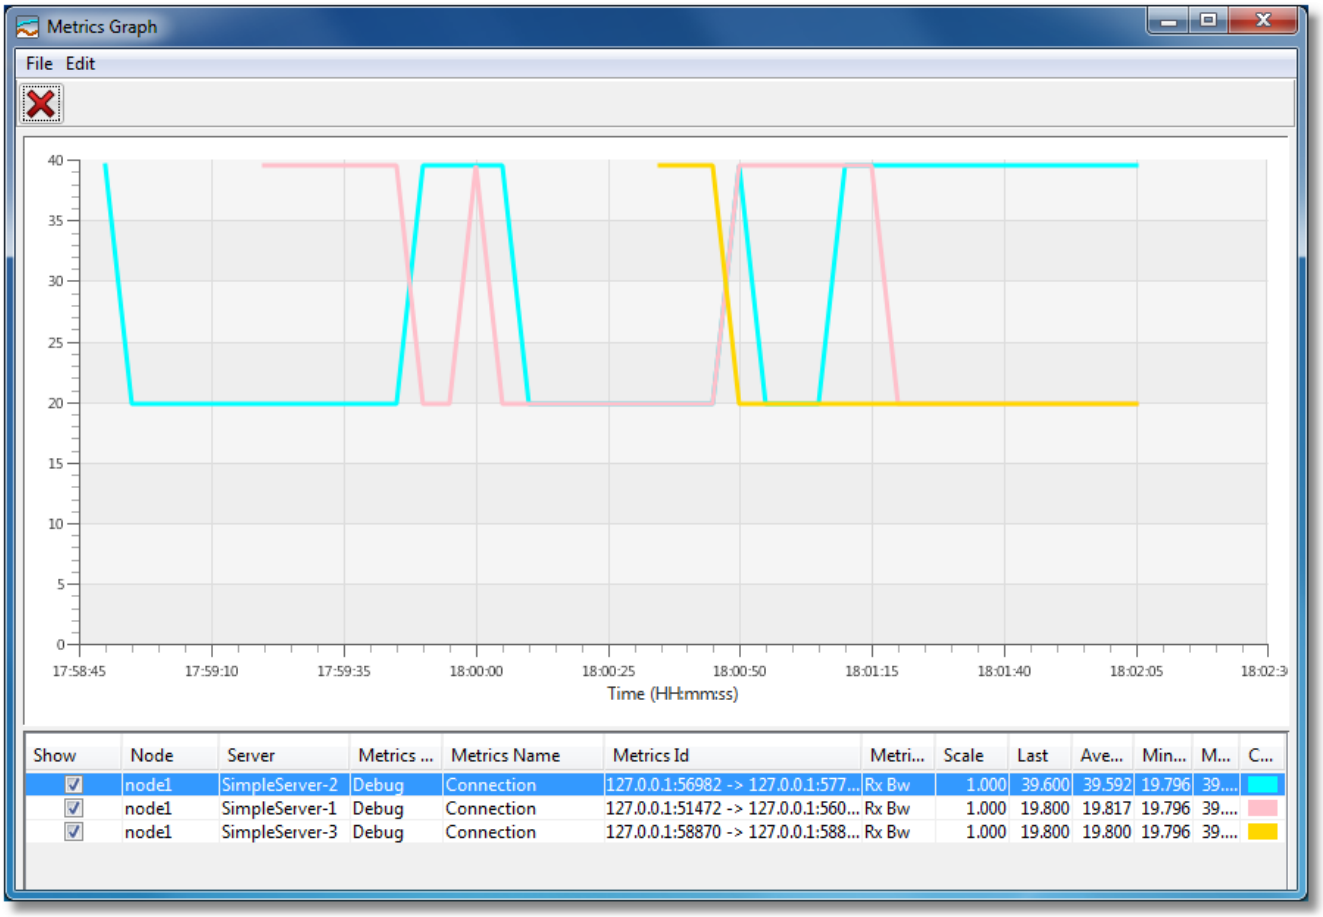

The horizontal axis corresponds to the time, while the vertical axis plots the corresponding metrics values.

> The Metrics Graph feature requires Java 7 and JavaFX.

On this page:

- Creating a New Metrics Graph
- Adding and Removing Metrics
- Adjusting Scale and Color
- Metrics Graph Preferences

Creating a New Metrics Graph

You can create a new Metrics Graph window with the `File > New > Metrics Graph` menu, or with a contextual menu over a metric in a Metrics View, for example:

Copyright 2016, ZeroC, Inc.

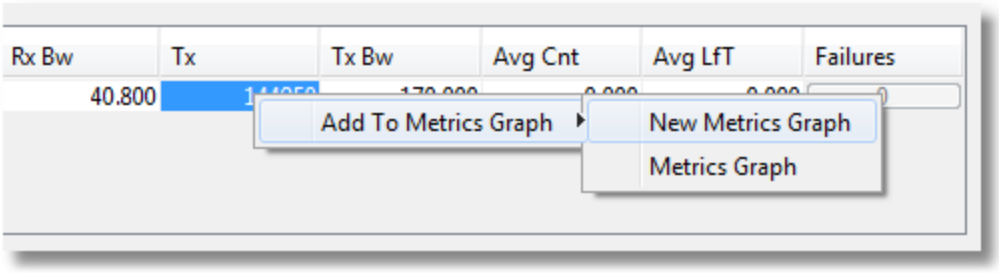

Adding and Removing Metrics

A metric is added to a Metrics Graph with drag & drop, or through a contextual menu over this metric.

To remove a metric from a Metrics Graph, select this metric and press Delete or the button.

Adjusting Scale and Color

You can display metrics with different scales on the same graph by adjusting the scale factor of each metric. The values plotted are the metric's values times the scale factor; the default scale factor is 1.0.

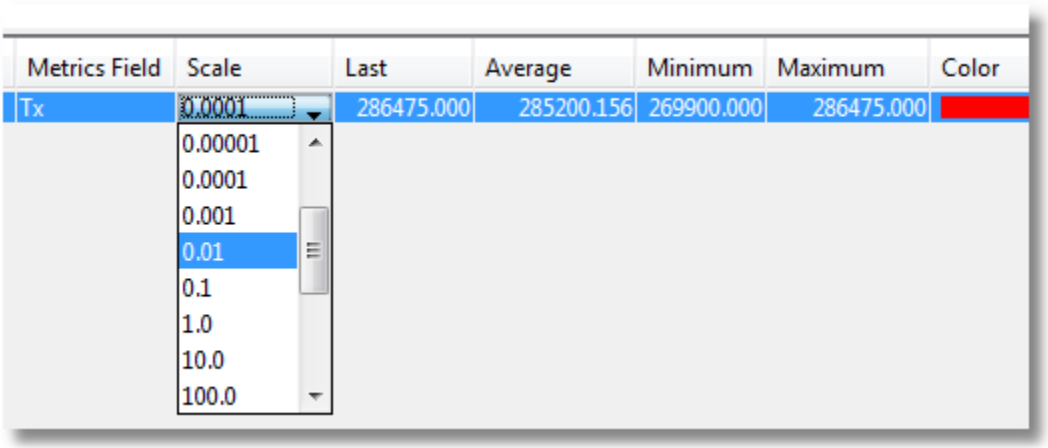

You can also change the color used to display a metric's line by clicking on the Color cell of the metric.

Metrics Graph Preferences

The `File > Preferences` menu of a Metrics Graph window opens a dialog that allows you to configure several properties of your Metrics Graph:

Copyright 2016, ZeroC, Inc.

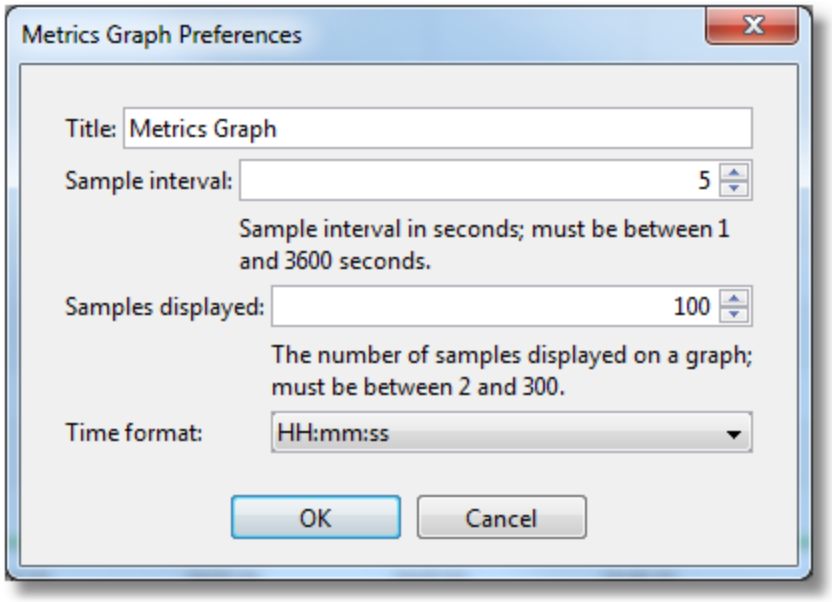

For each metric plotted in a Metrics Graph, IceGrid Admin retrieves the corresponding value every `Sample interval` seconds (by default every 5 seconds), and displays n values (n = 120 by default). `Sample interval` times `Samples displayed` correspond to the time period represented on a graph; by default, it is 5s * 120 = 10 minutes.

Copyright 2016, ZeroC, Inc.

Application Tabs

IceGrid definitions are organized in "applications", with typically one or a few applications deployed on a given IceGrid instance.

An Application tab allows you to:

- create a new application: describe servers, IceBox services, the nodes you want to run them, specify load-balancing policies, etc.
- browse and edit the definition of an existing, live application deployed in an IceGrid instance
- browse and edit the definition of an application stored in an IceGrid XML file.

Topics

- Editing and Saving IceGrid Descriptors
- Navigation within an Application Tab
- IceGrid Descriptors
  - Variables in IceGrid Descriptors
  - Application Descriptor
  - Node Descriptor
  - Server Descriptor
  - Service Descriptor
  - Adapter Descriptor
  - Database Environment Descriptor
  - Property Set Descriptor
  - Replica Group Descriptor
  - Server Template Descriptor
  - Service Template Descriptor

Copyright 2016, ZeroC, Inc.

Editing and Saving IceGrid Descriptors

On this page:

- Editing
- Copy & Paste
- Error Checking
- Saving
- Discarding Updates
- Concurrent Updates to the same IceGrid Registry

Editing

As soon as you make any update in a form, IceGrid Admin enables two buttons at the bottom of this form: Apply and Discard.

If you navigate to another node without pressing Apply or Discard, the default is Apply: your changes are applied to the in-memory representation of the application definition. However these changes are not stored to the IceGrid registry or XML file until you save the application definition (see below).

Editing a live application ( ) also disconnects this application from the IceGrid registry: updates made by other users are no longer propagated to the Application tab.

Copy & Paste

Most descriptor sub-trees can be copied and later pasted. Copies are always deep-copies: for example if you copy a node, all the servers on this code are copied, including all the the sub-elements of these servers (object adapters, database environment, services, etc.).

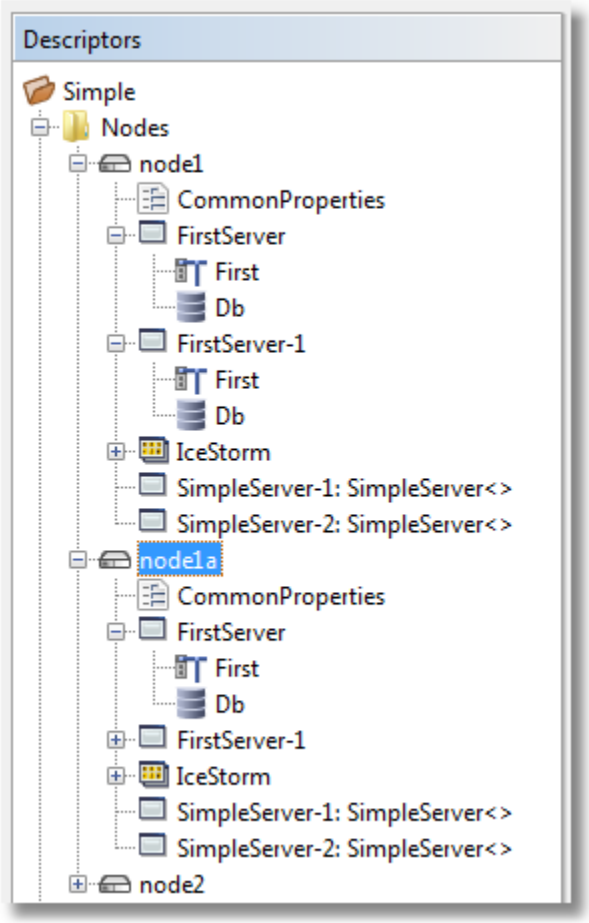

After pasting a sub-tree, you typically need to check and edit the new elements to avoid any duplicate server IDs, adapter IDs etc.

Error Checking

Copyright 2016, ZeroC, Inc.

IceGrid Admin performs very little error checking while you are working on an application definition. For example, you may temporarily keep several servers with the same ID, leave some parameters of a template instance unset, or use an undefined variable. All such errors are only detected when you save your application to an IceGrid registry.

There are nonetheless two types of constraints enforced by IceGrid Admin at all times:

- two descriptor nodes in an application tab display cannot have the same name.
- some fields (such as Path to Executable for a server) cannot be empty.

If you violate such a constraint, IceGrid Admin prevents you from applying your change.

Saving

You save an application definition to an IceGrid registry or an XML file with the menu item `File > Save`, `File > Save to File`, `File > Save to Registry (Servers may restart)`, `File > Save to Registry (No Server restart)` or with the corresponding toolbar buttons.

`Save` is equivalent to `Save to Registry (Servers may restart)` for a live application, and to `Save to File` for a file-bound application.

Saving an application to the IceGrid registry causes the registry to distribute any relevant changes to the affected nodes. In turn, each node applies those changes to its servers. If a server is running at the time of an update, the node may need to stop and restart it, depending on the changes that you make. Consequently, saving to the registry could cause a disruption in service to any clients that are actively using the affected servers.

The following application changes do *not* require a restart:

- Adding, modifying, or removing a server's configuration properties
- Adding new servers

All other changes will require a restart. IceGrid Admin provides two versions of the `Save to Registry` command, one that allows restarts and one that does not. To avoid accidentally causing any disruption in service, we recommend using the `No server restart` option first; this command will fail if any of your updates require a restart. At that point, you can decide whether to force the servers to restart using the other Save command.

If you change a server's configuration properties with `Save to Registry (No Server restart)`, IceGrid updates the stored properties of this server, and also the properties of your running server instance through its Properties Facet. If these properties are only read by the server at start-up, this may not have the desired effect. If you want to trigger a server restart even when only properties have changed, use `Save to Registry (Servers may restart)`.

Discarding Updates

You may discard all your updates by selecting `File > Discard Updates` or pressing the corresponding toolbar button. `Discard Updates` simply reloads the application from the IceGrid registry or its associated XML file.

Concurrent Updates to the same IceGrid Registry

If several administrators update the same application definition concurrently, the last save will silently overwrite previous (concurrent) updates.

To avoid this situation, you can acquire an exclusive write access to the IceGrid registry with `File > Acquire Exclusive Write Access`. After this exclusive write access is granted, any attempt by another session to save to the IceGrid registry will result in an error:

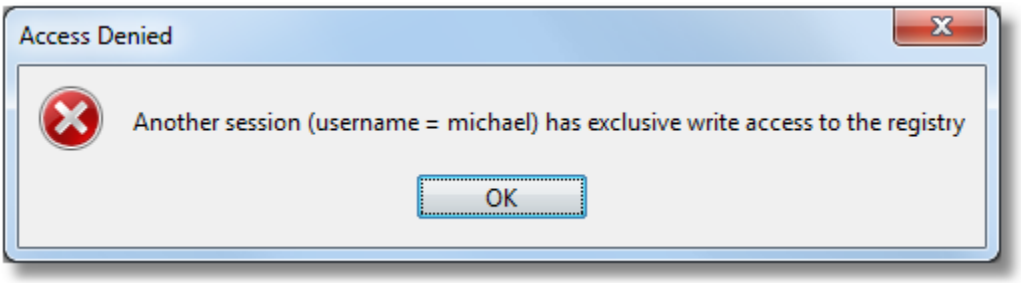

Copyright 2016, ZeroC, Inc.

Copyright 2016, ZeroC, Inc.

Navigation within an Application Tab

Each Application tab maintains a history of the nodes you have visited in this tab. You can navigate these nodes using the `View > Go Backto the Previous Node` and `View > Go to the Next Node` menu items, or with the equivalent toolbar buttons:

Also on some forms a green arrow button provides a link to another node (typically a template) in the same Application pane:

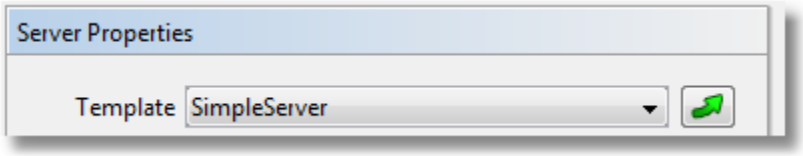

Copyright 2016, ZeroC, Inc.

IceGrid Descriptors

An application definition is described using one application descriptor with a number of nested sub-descriptors (and sub-sub descriptors). The descriptors correspond to Slice types in the IceGrid registry interfaces, and to XML element in an IceGrid XML representation.

IceGrid Admin maps these descriptors to nodes on a tree representation:

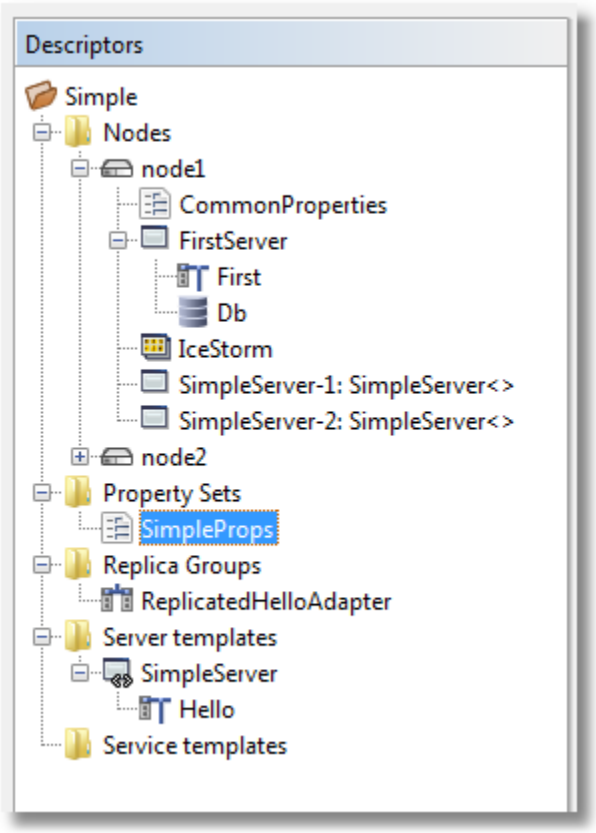

Topics

- Variables in IceGrid Descriptors
- Application Descriptor
- Node Descriptor
- Server Descriptor
- Service Descriptor
- Adapter Descriptor
- Database Environment Descriptor
- Property Set Descriptor
- Replica Group Descriptor
- Server Template Descriptor
- Service Template Descriptor

Copyright 2016, ZeroC, Inc.

Variables in IceGrid Descriptors

Variables allow you to define commonly-used information once and refer to them symbolically throughout your application descriptors.

On this page:

- Syntax
- Where are Variables Allowed?
- Escaping a Variable
- Pre-Defined Variables
- Variable Substitution in IceGrid Admin
- Scoping Rules
- Resolving a Reference
- Modifying a Variable

Syntax

Substitution for a variable or parameter VP is attempted whenever the symbol ${VP} is encountered, subject to the limitations and rules described below. Substitution is case-sensitive, and a fatal error occurs if VP is not defined when the application is saved to an IceGrid registry.

Where are Variables Allowed?

Substitution is performed in all string fields except the following:

- server and service template IDs (when defining a template or when referring to a template)
- variable names
- template parameter names
- node names
- application names

Escaping a Variable

You can prevent substitution by escaping a variable reference with an additional leading $ character. For example, in order to assign the literal string ${abc} to a variable, you would use $${abc} as this variable's value.

The extra $ symbol is only meaningful when immediately preceding a variable reference, therefore text such as US$$55 is not modified. Each occurrence of the characters $$ preceding a variable reference is replaced with a single $ character, and that character does not initiate a variable reference.

Pre-Defined Variables

IceGrid defines a set of read-only variables to hold information that may be of use to descriptors. The names of these variables are reserved and cannot be used as variable or parameter names. The table below describes the purpose of each variable and defines the context in which it is valid.

Name	Description
application	The name of the enclosing application.
application.distrib	The pathname of the enclosing application distribution directory, and an alias for ${node.datadir}/distrib/${application}.
node	The name of the enclosing node.
node.os	The name of the enclosing node operating system. On Unix, this is value is provided by uname. On Windows, the value is Windows.
node.hostname	The host name of the enclosing node.
node.release	The operation system release of the enclosing node. On Unix, this value is provided by uname. On Windows, the value is obtained from the OSVERSIONINFO data structure.
node.version	The operation system version of the enclosing node. On Unix, this value is provided by uname. On Windows, the value represents the current service pack level.
node.machine	The machine hardware name of the enclosing node. On Unix, this value is provided by uname. On Windows, the value is x86 or x64.

Copyright 2016, ZeroC, Inc.

node.datadir	The absolute pathname of the enclosing node data directory.
server	The ID of the enclosing server.
server.distrib	The pathname of the enclosing server distribution directory, and an alias for ${node.datadir}/servers/${server}/distrib.
service	The name of the enclosing service.
session.id	The client session identifier. For sessions created with a user name and password, the value is the user ID; for sessions created from a secure connection, the value is the distinguished name associated with the connection.

The availability of a variable is easily determined in some cases, but may not be readily apparent in others. For example, you can use the ${node} variable in a property value within a server template definition, because variables in the body of a server template are evaluated when the server template is instantiated on a specific node.

Variable Substitution in IceGrid Admin

In a number of panes, you can substitute variables and template parameters by their respective value. Use `View > Show Variables` and `View > Substitute Variables` or the corresponding toolbar toggle buttons:

When variable-substitution is enabled, the descriptors are displayed read-only.

Scoping Rules

Descriptors may only define variables at the application and node levels. Each node introduces a new scope, such that defining a variable at the node level overrides (but does not modify) the value of an application variable with the same name. Similarly, a template parameter overrides the value of a variable with the same name in an enclosing scope. A descriptor may refer to a variable defined in any enclosing scope, but its value is determined by the nearest scope. The diagram below illustrates these concepts:

Copyright 2016, ZeroC, Inc.

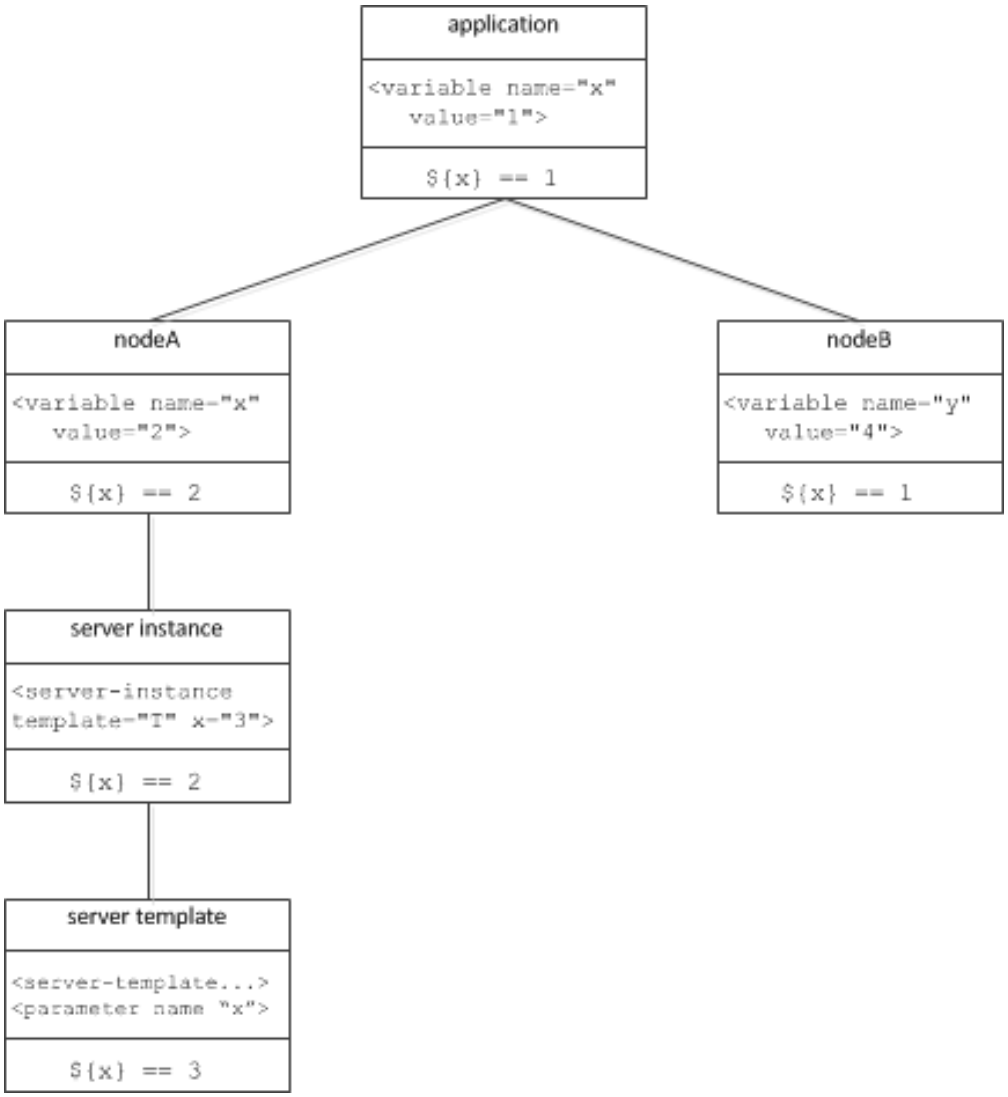

In this diagram, the variable x is defined at the application level with the value 1. In node A, x is overridden with the value 2, whereas x remains unchanged in node B. Within the context of node A, x continues to have the value 2 in a server instance definition. However, when x is used as the name of a template parameter, the node's definition of x is overridden and x has the value 3 in the template's scope.

Resolving a Reference

To resolve a variable reference ${var}, IceGrid searches for a definition of var using the following order of precedence:

- Pre-defined variables
- Template parameters, if applicable
- Node variables, if applicable
- Application variables

After the initial substitution, any remaining references are resolved recursively using the following order of precedence:

- Pre-defined variables
- Node variables, if applicable
- Application variables
- Template Parameters

Template parameters are not visible in nested template instances. This situation can only occur when an IceBox server template instantiates a service template.

Modifying a Variable

A variable definition can be overridden in an inner scope, but the inner definition does not modify the outer variable.
See Also

Copyright 2016, ZeroC, Inc.

- Using Descriptor Variables and Parameters

Copyright 2016, ZeroC, Inc.

Application Descriptor

On this page:

- Creating a new Application
- Properties
- Children

Creating a new Application

You can create a new application using `File > New Application`: this opens a new Application tab, with an empty application definition.

When you are connected to an IceGrid registry, you can also use `File > New Application with Default Templates from Registry`. This also opens new Application tab with a brand new application definition. This new application contains a copy of all the templates definitions contained in the IceGrid registry default template file. See `IceGrid.Registry.DefaultTemplates`.

Properties

The Application Properties panel offers the following fields:

- **Name**
  The name of the application. This field is not editable for live applications.

- **Description**
  A free-text description of this application.

- **Variables**
  This table shows application-level IceGrid variables.

- **IcePatch2 Proxy**
  A stringified or well-known proxy for the IcePatch2 server than contains this application's distribution. Possible values:

Value	Description
"None selected" (default) or blank	no distribution defined
`${application}.IcePatch2/server`	the well-known proxy of an IcePatch2 server deployed in this application using the IcePatch2 default server-template (with a default instance-name parameter value).
stringified proxy	your own stringified or well-known proxy

- **Directories**
  List of directories included in the application distribution. When blank, the entire IcePatch2 server repository is used as the distribution.

Children

An application node can have five types of children:

- Node
- Property Set
- Replica Group
- Server Template
- Service Template

Copyright 2016, ZeroC, Inc.

Node Descriptor

A node represents an IceGrid node that starts and monitors your servers. Several applications can be deployed on the same node; however a node descriptor describes only the servers defined by the enclosing application.

On this page:

- Properties
- Children

Properties

The Node Properties panel offers the following fields:

- **Name**
  The name of the node. This name must match the IceGrid.Node.Name configuration property of the node process.

- **Description**
  A free-text description of this node.

- **Variables**
  This table shows node-level IceGrid variables.

- **Load Factor**
  A floating point value used to compare different nodes when making a load-balancing decision based on load-average (for Linux and Unix) or CPU utilization (for Windows). Leaving this value blank is equivalent to the default: 1.0 on Linux and Unix, and 1.0 divided by the number of CPUs on Windows.

Children

A node can have two types of children:

- Server
- Property Set

Copyright 2016, ZeroC, Inc.

Server Descriptor

A server represents an Ice server deployed on a node as part of an application. IceGrid supports three kinds of servers:

- **Plain Server**
  A normal Ice server, with typically a single Ice communicator. IceGrid generates a single Ice configuration file for this server.

- **IceBox Server**
  An IceBox server, with usually a number of IceBox services deployed in it. IceGrid generates an Ice configuration file for the IceBox server itself, and also a separate Ice configuration for each IceBox service.

- **Server Instance**
  A server (either plain server or IceBox server) defined using a server template.

On this page:

- Plain Server
  - Properties
  - Children
- IceBox Server
  - Properties
  - Children
- Server Instance
  - Properties
  - Children

Plain Server

Properties

The Server Properties panel offers the following fields:

- **Server ID**
  The ID of the server; corresponds to the `Ice.Admin.ServerId` property. Each server must have a unique ID within an IceGrid deployment.

- **Description**
  A free-text description of this server.

- **Property Sets**
  List of property-set IDs; you refer to a property set to "include" all its properties in the server.

- **Properties**
  Ice properties private to this server.

- **Log Files**
  This table can be used to declare a number of log files used by this server. Path is the path to the log file (a relative path is relative to the IceGrid node working directory); when Property is set, IceGrid generates a property with this name and the log file path as value.
  You declare log files to be able to conveniently retrieve them using IceGrid Admin (in the Live Deployment tab) or with the `icegrid admin` command-line utility.

- **Path to Executable**
  Path to the server's executable; cannot be blank. A relative path is relative to the IceGrid node working directory.

- **Ice Version**
  The Ice version of this server. If you don't provide a value, IceGrid assumes it's the same version as the IceGrid registry.

- **Working Directory**
  The working directory for the server when started by the IceGrid node.

- **Command Arguments**
  The command-line arguments given to the server when started by the IceGrid node.

- **Run as**
  On Linux and Unix, when IceGrid node is running as root, it is possible to run the server under any username. Enter the desired username in this field.

Copyright 2016, ZeroC, Inc.

When not set (the default), the server runs as the same user as the IceGrid node, except when IceGrid node runs as root. In this case, the server runs as nobody.

- **Environment Variables**
  The environment variables for the server when started by the IceGrid node. These variables are in addition to variables defined in the IceGrid node own environment.

- **Activation Mode**
  The server's activation mode. Must be one of:

Activation Mode	Description
always	IceGrid keeps this server running all the time
manual	This server is started "manually", using IceGrid Admin or the `icegridadmin` command-line utility
on-demand	IceGrid starts this server when it resolves the object-adapter ID of an object adapter defined in this server
session	IceGrid starts a separate instance of this server for each IceGrid session that allocates this server

The Activation Mode can also be a variable or a combination of variables that resolves to one of the values above.

- **Activation Timeout**
  When activating a server, IceGrid gives timeout seconds to object adapters with server lifetime to register their endpoints with the IceGrid registry. During this time, lookup for the corresponding adapter IDs are delayed.
  If not set or set to 0, the IceGrid node uses the value of its `IceGrid.Node.WaitTime` property.

- **Deactivation Timeout**
  When deactivating a server, IceGrid gives timeout seconds to the server to exit gracefully. After this timeout, the server process is killed.
  If not set or set to 0, the IceGrid node uses the value of its `IceGrid.Node.WaitTime` property.

- **Allocatable**
  Specifies whether the server can be allocated. A server is allocated implicitly when one of its allocatable objects is allocated. This checkbox is ignored if the server activation mode is session; a server with this activation mode is always allocatable. Default: false.

- **Depends on the application distribution**
  A server that depends on the application distribution is stopped and disabled before any application distribution update, and re-enabled after such update. Default: true.

- **IcePatch2 Proxy**
  A stringified (or well-known) proxy for the IcePatch2 server than contains this server's private distribution. Possible values:

Value	Description
"None selected" (default) or blank	no private distribution defined
`${application}.IcePatch2/server`	the well-known proxy of an IcePatch2 server deployed in this application using the IcePatch2 default server-template (with a default instance-name parameter value).
stringified proxy	your own stringified or well-known proxy

- **Directories**
  List of directories included in the application distribution. When blank, the entire IcePatch2 server repository is used as the distribution.

Children

A plain server can have two types of children:

- Adapter
- Database Environment

IceBox Server

Copyright 2016, ZeroC, Inc.

Properties

The Properties panel for an IceBox server is identical to the Properties panel for a Plain Server.

When you create a new IceBox server, some properties are created automatically:

Property Name	Property Value
IceBox.InstanceName	${server}
Ice.Admin.Endpoints	tcp -h 127.0.0.1

The `Ice.Admin.Endpoints` setting enables the Admin object in the main communicator of this IceBox server.

The Path to Executable is typically `icebox`, `iceboxd` (for a C++ IceBox), `java` (for a Java IceBox) or `iceboxcs` (for a .NET IceBox). In Java, the Command Arguments should include the IceBox container class name, `IceBox.Server`.

Children

An IceBox server can have only one type of children, Service.

Server Instance

A server instance is a server created from a server template; it may be a plain server or an IceBox server.

Properties

The Server Instance Properties panel offers the following fields:

- **Template**
  The name of the Server Template.
- **Parameters**
  Use this table to assign values to the template parameters defined in the server template.
- **Property Sets**
  List of property-set IDs; you refer to a property set to "include" all its properties in this server-instance.
- **Properties**
  Ice properties private to this server. Overall, the properties of the server instance are a combination of properties defined in the server template (including its own property sets references) augmented and possibly overridden by properties defined in the server instance.

Children

An instance of an IceBox server template can have one type of children, Property Set.

An instance of a plain server template cannot have any child.

Copyright 2016, ZeroC, Inc.

Service Descriptor

A service represents an IceBox service deployed on an IceBox server. IceGrid supports two kinds of services:

- **Plain Service**
  A service defined directly with a service descriptor

- **Service Instance**
  A service defined using a service template

On this page:

- Load Order
- Plain Service
  - Properties
  - Children
- Service Instance
  - Properties

Load Order

The load-order of services within an IceBox server is determined by their display order in IceGrid Admin. You can reorder services using `Edit > Move Up` or `Edit > Move Down` (also available from each service's contextual menu).

Plain Service

Properties

The Service Properties panel offers the following fields:

- **Service Name**
  The name of the service. Must be unique within the enclosing IceBox server.

- **Description**
  A free-text description of this service.

- **Property Sets**
  List of property-set IDs; you refer to a property set to "include" all its properties in the service.

- **Properties**
  Ice properties private to this service.

- **Log Files**
  This table can be used to declare a number of log files used by this service. Path is the path to the log file (a relative path is relative to the IceGrid node working directory); when Property is set, IceGrid generates a property with this name and the log file path as value.
  You declare log files to be able to conveniently retrieve them using IceGrid Admin (in the Live Deployment pane) or with the `icegridadmin` command-line utility.

- **Entry Point**
  The service's entry point, for example `IceStormService,35:createIceStorm`.
  IceGrid will append `--Ice.Config=`*`path-to-service-config-file`* to this entry point when it generates the `IceBox.Service.`*`name`* property in the enclosing IceBox server config file.

Children

A plain service can have two types of children:

- Adapter
- Database Environment

Service Instance

Properties

Copyright 2016, ZeroC, Inc.

The Service Instance Properties panel offers the following fields:

- **Template**
  The name of the Service Template.

- **Parameters**
  Use this table to give values to the template parameters defined in the service template.

- **Property Sets**
  List of property-set IDs; you refer to a property set to "include" all its properties in the service.

- **Properties**
  Ice properties private to this service.

Copyright 2016, ZeroC, Inc.

Adapter Descriptor

Each indirect object adapter registered with IceGrid requires its own Adapter descriptor. If you need to specify a direct adapter, simply create a number of Ice properties in your server.

Properties

The Adapter Properties panel offers the following fields:

- **Adapter Name**
  The name of the object adapter. This name must be unique within the enclosing Ice communicator, and is only used to lookup adapter properties.
- **Description**
  A free-text description of this object adapter.
- **Adapter ID**
  The ID of the object adapter. This ID must be unique within an IceGrid deployment. Default value: ${server}.*adapter-name*.
- **Replica Group**
  The ID of this adapter's Replica Group. By default, an adapter does not belong to any replica group.
- **Priority**
  The adapter priority in its Replica Group. The default priority is 0.
- **Endpoints**
  The configured endpoints for this object adapter. Corresponds to the *`adapter-name`*`.Endpoints` property. Default: default, which means listen using the default protocol (tcp by default) on an OS assigned port, on all interfaces.
- **Published Endpoints**
  The configured published endpoints for this object adapter. Corresponds to the *`adapter-name`*`.PublishedEndpoints` property. Default: actual endpoints (computed at run time), derived by Ice from the Endpoints field above.
- **Proxy Options**
  The default Proxy Options for well-known objects defined within this object adapter, and the Proxy Options for proxies created by this object adapter. Corresponds to the *`adapter-name`*`.ProxyEndpoints` property. Default: empty string.
- **Register Process**
  This setting is ignored for servers running Ice version 3.3 or greater.
  When checked, this adapter will create and register an `Ice::Process` object with IceGrid during activation. Such object is used to cleanly shutdown the server from IceGrid. Each server should register one, and only one such `Ice::Process` object.
- **Server Lifetime**
  When checked, IceGrid expects this adapter to register its endpoints during server startup and unregister them during server shutdown. See also Activation Timeout in Server Properties. Default: true (checked).
- **Well-known Objects**
  A table of well-known objects defined by this adapter. When Property is set, IceGrid generates a property with this name and with Identity as value.
- **Allocatable Objects**
  A table of allocatable objects defined by this adapter. When Property is set, IceGrid generates a property with this name and with Identity as value.

Copyright 2016, ZeroC, Inc.

Database Environment Descriptor

You can describe the Freeze database environment(s) used by your server or service with Database Environment Descriptors. Each descriptor describes a single Freeze database environment.

Properties

The Database Environment Properties panel offers the following fields:

- **Name**
  The name of the Freeze Database Environment. This name must be unique within the enclosing Ice communicator.

- **Description**
  A free-text description of this database environment.

- **DB Home**
  The path of the home directory of the Berkeley DB environment. The default is "Created by the IceGrid Node": when selected, IceGrid node creates automatically a directory in its local server tree. Otherwise, if you enter your own value, the path must exist: IceGrid node will not create this directory if it does not exist.

- **Properties**
  This table shows a list of Berkeley DB properties for this database environment. These are *not* Ice properties. IceGrid generates a `DB_CONFIG` file in the DB Home directory with these settings.

Copyright 2016, ZeroC, Inc.

Property Set Descriptor

A Property Set defines a set of Ice properties. IceGrid Admin supports two kinds of Property Sets:

- **Named Property Set**
  A property set defined within an application or within a node. Server and service definitions refer to such property sets to "include" the corresponding properties.

- **Service-Instance Property Set**
  A property set defined as a child of an IceBox server instance. Such a property set provides properties to a service instance within a concrete IceBox server.

Properties

The Property Set Properties panel offers the following fields:

- **ID** (Named Property Set only)
  The ID of a named property set.

- **Service Name** (Service Instance property set)
  The service name.

- **Property Sets**
  List of Property Set IDs. The corresponding property sets are "included" in this property set.

- **Properties**
  Ice properties private to this Property Set

Copyright 2016, ZeroC, Inc.

Replica Group Descriptor

A replica group represents an abstract grouping of identical, or very similar, object adapters. An object adapter joins this group by setting this replica group ID in its Replica Group field.

Properties

The Replica Properties panel offers the following fields:

- **Replica Group ID**
  The ID of this replica group. Must be unique within an IceGrid deployment.
- **Description**
  A free-text description of this replica group.
- **Proxy Options**
  The default Proxy Options for well-known objects defined within this replica group.
- **Well-known Objects**
  A table of well-known objects defined by this replica group.
- **Load Balancing Policy**
  The load-balancing policy used by IceGrid when resolving the Replica Group ID in proxies. Must be one of:

Policy	Description
Adaptive	Returns the endpoints of the object adapters whose nodes report the lowest load-average (on Linux/Unix) or CPU utilization (on Windows). The actual load-average or CPU utilization value of each node is multiplied by the node's load-factor for this comparison. See Load Sample below, and Node.
Ordered	Returns the endpoints of the object adapters with the highest priority. See Adapter.
Random	Returns the endpoints of object adapters selected at random. This is the default policy.
Round-robin	IceGrid puts all the object adapters in a list, and for each new resolution, it returns the endpoints from the next *How many Adapters* items on this list.

- **How many Adapters**
  Specify the number of object adapters selected by IceGrid for each resolution. 1 is a common value. When set to 0 (the default), IceGrid returns the endpoints of all object adapters in this replica group.
- **Load Sample**
  Available only with the Adaptive load-balancing policy. Specify the load-average or CPU utilization sample to use when comparing nodes: it can be 1 (default), 5 or 15 minutes.

Copyright 2016, ZeroC, Inc.

Server Template Descriptor

A server template is used to capture the common definitions of several similar or identical servers.

On this page:

- Properties
- Children

Properties

The Server Template Properties panel offers the following fields:

- **Template ID**
  The ID the template. Must be unique within the application.

- **Parameters**
  The list of parameters for this template. Each parameter can have an optional default value.

- **Plain Server Properties**
  The remaining fields are the Plain Server fields.

Children

A plain server template can have two types of children:

- Adapter
- Database Environment

An IceBox server template can have only one kind of children, Service.

Copyright 2016, ZeroC, Inc.

Service Template Descriptor

A service template is used to capture the common definitions of several similar or identical services.

On this page:

- Properties
- Children

Properties

The Service Template Properties panel offers the following fields:

- **Template ID**
  The ID the template. Must be unique within the application.

- **Parameters**
  The list of parameters for this template. Each parameter can have an optional default value.

- **Plain Service Properties**
  The remaining fields are the Plain Service fields.

Children

A service template can have two types of children:

- Adapter
- Database Environment

Copyright 2016, ZeroC, Inc.

IceGrid Server Activation

On this page:

- Server Activation Modes
- Server Activation in Detail
- Requirements for Server Activation
- Efficiency Considerations for Server Activation
- Activating Servers with Specific User IDs
- Automating Endpoint Registration

Server Activation Modes

You can choose among four activation modes for servers deployed and managed by an IceGrid node:

- Manual
  You must start the server explicitly via the IceGrid GUI or `icegridadmin`, or programmatically via the `IceGrid::Admin` interface.
- Always
  IceGrid activates the server when its node starts. If the server stops, IceGrid automatically reactivates it.
- On demand
  IceGrid activates the server when a client invokes an operation on an object in the server.
- Session
  This mode also provides on-demand activation but requires the server to be allocated by a session.

Server Activation in Detail

On-demand server activation is a valuable feature of distributed computing architectures for a number of reasons:

- It minimizes application startup times by avoiding the need to pre-start all servers.
- It allows administrators to use their computing resources more efficiently because only those servers that are actually needed are running.
- It provides more reliability in the case of some server failure scenarios, e.g., the server is reactivated after a failure and may still be capable of providing some services to clients until the failure is resolved.
- It allows remote activation and deactivation.

On-demand activation occurs when an Ice client requests the endpoints of one of the server's object adapters via a locate request. If the server is not active at the time the client issues the request, the node activates the server and waits for the target object adapter to register its endpoints. Once the object adapter endpoints are registered, the registry returns the endpoint information back to the client. This sequence ensures that the client receives the endpoint information *after* the server is ready to receive requests.

Requirements for Server Activation

In order to use on-demand activation for an object adapter, the adapter must have an identifier and be entered in the IceGrid registry.

When using session activation mode, IceGrid requires that the server be allocated; on-demand activation fails for servers that have not been allocated.

The session activation mode recognizes an additional reserved variable in the server descriptor, `${session.id}`. The value of this variable is the user ID or, for SSL sessions, the distinguished name associated with the session.

Efficiency Considerations for Server Activation

Once a server is activated, it remains running indefinitely (unless it uses the session activation mode). A node deactivates a server only when explicitly requested to do so. As a result, server processes tend to accumulate on the node's host.

One of the advantages of on-demand activation is the ability to manage computing resources more efficiently. Of course there are many aspects to this, but Ice makes one technique particularly simple: servers can be configured to terminate gracefully after they have been idle for a certain amount of time.

A typical scenario involves a server that is activated on demand, used for a while by one or more clients, and then terminated automatically when no requests have been made for a configurable number of seconds. All that is necessary is setting the server's configuration property `Ice.ServerIdleTime` to the desired idle time.

For a server activated in session activation mode, IceGrid deactivates the server when the session releases the server or when the session is destroyed.

Copyright 2016, ZeroC, Inc.

Activating Servers with Specific User IDs

On Unix platforms you can activate server processes with specific effective user IDs, provided that the IceGrid node is running as root. If the IceGrid node does not run as root, servers are always activated with the effective user ID of the IceGrid node process. (The same is true for Windows — servers always run with the same user ID as the IceGrid node process.)

For the remainder of this section, we assume that the node runs as root on a Unix machine.

The `user` attribute of the server descriptor specifies the user ID for a server. If this attribute is not specified and the activation mode is not `session`, the default value is `nobody`. Otherwise, the default value is `${session.id}` if the activation mode is `session`.

Since individual users often have different account names and user IDs on different machines, IceGrid provides a mechanism to map the value of the `user` attribute in the server descriptor to a user account. To do this, you must configure the node to use a user account mapper object. This object must implement the `IceGrid::UserAccountMapper` interface:

Slice

```
exception UserAccountNotFoundException {};

interface UserAccountMapper {
    string getUserAccount(string user)
            throws UserAccountNotFoundException;
};
```

The IceGrid node invokes `getUserAccount` and passes the value of the server descriptor's `user` attribute. The return value is the name of the user account.

IceGrid provides a built-in file-based user account mapper that you can configure for the node and the registry. The file contains any number of user-account-ID pairs. Each pair appears on a separate line, with white space separating the user account from the identifier. For example, the file shown below contains two entries that map two distinguished names to the user account `lisa`:

```
lisa O=ZeroC\\, Inc., OU=Ice, CN=Lisa
lisa O=ZeroC\\, Inc., OU=Ice, CN=Lisa S.
```

The distinguished names must be unique. If the same distinguished name appears several times in a file, the last entry is used.

You can specify the path of the user account file with the `IceGrid.Registry.UserAccounts` property for the registry and the `IceGrid.Node.UserAccounts` property for a node.

To configure an IceGrid node to use the IceGrid registry file-based user account mapper, you need to set the `IceGrid.Node.UserAccountMapper` property to the well-known proxy `IceGrid/RegistryUserAccountMapper`. Alternatively, you can set this property to the proxy of your own user account mapper object. Note that if this property is set, the node ignores the setting of `IceGrid.Node.UserAccounts`.

Automating Endpoint Registration

Servers must be properly configured to enable automatic endpoint registration. It should be noted however that IceGrid simplifies the configuration process in two ways:

- The IceGrid deployment facility automates the creation of a configuration file for the server, including the definition of object adapter identifiers and endpoints.
- A server that is activated automatically by an IceGrid node does not need to explicitly configure a proxy for the locator because the IceGrid node defines it in the server's configuration file.

See Also

- Getting Started with IceGrid
- IceGrid Architecture

Copyright 2016, ZeroC, Inc.

- Resource Allocation using IceGrid Sessions
- Server Descriptor Element
- Locator Configuration for a Server
- Using IceGrid Deployment
- IceGrid.*

Copyright 2016, ZeroC, Inc.

IceGrid Troubleshooting

On this page:

- Troubleshooting Activation Failures
- Troubleshooting Proxy Failures
- Troubleshooting Server Failures
- Disabling Faulty Servers

Troubleshooting Activation Failures

Server activation failure is usually indicated by the receipt of a `NoEndpointException`. This can happen for a number of reasons, but the most likely cause is an incorrect configuration. For example, an IceGrid node may fail to activate a server because the server's executable file, shared libraries, or classes could not be found. There are several steps you can take in this case:

1. Enable activation tracing in the node by setting the configuration property `IceGrid.Node.Trace.Activator=3`.
2. Examine the tracing output and verify the server's command line and working directory are correct.
3. Relative pathnames specified in a command line may not be correct relative to the node's current working directory. Either replace relative pathnames with absolute pathnames, or restart the node in the proper working directory.
4. Verify that the server is configured with the correct `PATH` or `LD_LIBRARY_PATH` settings for its shared libraries. For a Java server, its `CLASSPATH` may also require changes.

Another cause of activation failure is a server fault during startup. After you have confirmed that the node successfully spawns the server process using the steps above, you should then check for signs of a server fault (e.g., on Unix, look for a `core` file in the node's current working directory).

Troubleshooting Proxy Failures

A client may receive `Ice::NotRegisteredException` if binding fails for an indirect proxy. This exception indicates that the proxy's object identity or object adapter is not known by the IceGrid registry. The following steps may help you discover the cause of the exception:

1. Use `icegridadmin` to verify that the object identity or object adapter identifier is actually registered, and that it matches what is used by the proxy:

```
>>> adapter list
...
>>> object find ::Hello
...
```

2. If the problem persists, review your configuration to ensure that the locator proxy used by the client matches the registry's client endpoints, and that those endpoints are accessible to the client (i.e., are not blocked by a firewall).
3. Finally, enable locator tracing in the client by setting the configuration property `Ice.Trace.Locator=2`, then run the client again to see if any log messages are emitted that may indicate the problem.

Troubleshooting Server Failures

Diagnosing a server failure can be difficult, especially when servers are activated automatically on remote hosts. Here are a few suggestions:

1. If the server is running on a Unix host, check the current working directory of the IceGrid node process for signs of a server failure, such as a `core` file.
2. Judicious use of tracing can help to narrow the search. For example, if the failure occurs as a result of an operation invocation, enable protocol tracing in the Ice run time by setting the configuration property `Ice.Trace.Protocol=1` to discover the object identity and operation name of all requests.
   Of course, the default log output channels (standard out and standard error) will probably be lost if the server is activated automatically, so either start the server manually (see below) or redirect the log output. You can also use the `Ice::Logger` interface to emit your own trace messages.
3. Run the server in a debugger; a server configured for automatic activation can also be started manually if necessary. However, since the IceGrid node did not activate the server, it cannot monitor the server process and therefore will not know when the server terminates. This will prevent subsequent activation unless you clean up the IceGrid state when you have finished debugging and terminated the server. You can do this by starting the server using `icegridadmin`:

Copyright 2016, ZeroC, Inc.

```
>>> server start TheServer
```

This will cause the node to activate (and therefore monitor) the server process. If you do not want to leave the server running, you can stop it with the `server stop` command.

4. After the server is activated and is in a quiescent state, attach your debugger to the running server process. This avoids the issues associated with starting the server manually (as described in the previous step), but does not provide as much flexibility in customizing the server's startup environment.

Another cause for a server to fail to activate correctly is if there is a mismatch in the adapter identifiers used by the server for its adapters, and the adapter identifiers specified in the server's deployment descriptor. After starting a server process, the node waits for the server to activate all of its object adapters and report them as ready; if the server does not do this, the node reports a failure once a timeout expires. The timeout is controlled by the setting of the property `IceGrid.Node.WaitTime`. (The default value is 60 seconds.)

You can check the status of each of a server's adapters using `icegridadmin` or the GUI tool. While the node waits for an adapter to be activated by the server, it reports the status of the adapter as "activating". If you experience timeouts before each adapter's status changes to "active", the most likely cause is that the deployment descriptor for the server either mentions more object adapters than are actually created by the server, or that the server uses an identifier for one or more adapters that does not match the corresponding identifier in the deployment descriptor.

Disabling Faulty Servers

You may find it necessary to disable a server that terminates in an error condition. For example, on a Unix platform each server failure might result in the creation of a new (and potentially quite large) core file. This problem is exacerbated when the server is used frequently, in which case repeated cycles of activation and failure can consume a great deal of disk space and threaten the viability of the application as a whole.

As a defensive measure, you can configure an IceGrid node to disable these servers automatically using the `IceGrid.Node.DisableOnFailure` property. In the disabled state, a server cannot be activated on demand. The default value of the property is zero, meaning the node does not disable a server that terminates improperly. A positive value causes the node to temporarily disable a faulty server, with the value representing the number of seconds the server should remain disabled. If the property has a negative value, the server is disabled indefinitely, or until the server is explicitly enabled or started via an administrative action.

You can also manually disable a server at any time using an administrative tool. A manually disabled server remains disabled indefinitely until an administrator enables or starts it. Disabling an *active* server has no effect on the server process; the server is unaware of the change to its status and continues to service requests from connected clients as usual. However, as of Ice 3.5, disabling a server does prevent IceGrid from including the endpoints of the server's object adapters in any subsequent locate requests, and it excludes those object adapters from any replica groups in which they might participate.

Typically, the ultimate goal of disabling a server is to gracefully migrate clients from the faulty server to ones that are behaving correctly. For a client that starts after the server is disabled, migration occurs immediately: the Ice run time in the client issues a locate request on the first proxy invocation, and the result returned by IceGrid excludes any endpoints from the disabled server. For a client that is active at the time the server is disabled, it does not migrate to a different server until its Ice run time issues a new locate request and obtains a new set of endpoints for the target object. The timing of this new locate request depends on several factors:

- If the client has an existing connection to the server, that connection will remain open and active as determined by the configuration settings of the client and server. For example, active connection management could eventually cause the connection to be closed automatically from either end due to inactivity. A subsequent proxy invocation may result in a new locate request.
- The client's use of connection caching also plays an important role. If the client's proxy is configured to cache connections, the client may continue to use the disabled server indefinitely. Connection caching should be disabled to ensure that migration eventually takes place.
- The Ice run time in the client also caches the results of locate requests. Although by default these results do not expire, setting a cache timeout allows you to specify how frequently the Ice run time issues new locate requests. Consider this code:

C++

```
proxy = proxy->ice_connectionCached(false)->ice_locatorCacheTimeout(20);
```

By disabling connection caching and setting a locator cache timeout, we can ensure that migration occurs within twenty seconds for invocations on this proxy.

See Also

Copyright 2016, ZeroC, Inc.

- IceGrid Server Activation
- Locator Semantics for Clients
- icegridadmin Command Line Tool

Copyright 2016, ZeroC, Inc.

IceGrid Database Utility

The `icegriddb` utility is a command-line tool for importing and exporting an IceGrid registry database. This tool is especially useful when migrating to Ice 3.7, where IceGrid uses a different database implementation and storage format. Using `icegriddb`, you can export your existing IceGrid 3.6 registry database and then import the data using the Ice 3.7 version of `icegriddb`.

On this page:

- Usage
- Exporting an IceGrid Database
- Importing an IceGrid Database

Usage

The IceGrid Database utility supports the following command-line options:

```
Usage: icegriddb <options>
Options:
-h, --help            Show this message.
-v, --version         Display version.
--import FILE         Import database from FILE.
--export FILE         Export database to FILE.
--dbhome DIR          The database directory.
-d, --debug           Print debug messages.
```

Exporting an IceGrid Database

To export an IceGrid registry database, use the `--export` option to specify the output file and the `--dbhome` option to specify the path name of the registry's database environment. To discover the location of your database environment, review the registry's configuration and look for the setting of `IceGrid.Registry.Data`. For example, the IceGrid sample programs typically use this setting:

```
IceGrid.Registry.Data=db/registry
```

Run the following command to export the database:

```
$ icegriddb --export registry.ixp --dbhome db/registry
```

The utility will fail with an error message if the given path does not contain a registry database.

Importing an IceGrid Database

To import an IceGrid registry database, use the `--import` option to specify the input file and the `--dbhome` option to specify the path name of the registry's database environment. For example, use the following command to import a database into the `dbNew/registry` directory from a file named `registry.ixp`:

```
$ icegriddb --import registry.ixp --dbhome dbNew/registry
```

The target directory must be empty.

Copyright 2016, ZeroC, Inc.

IcePatch2

IcePatch2 is the Ice solution for secure replication of a directory tree.

IcePatch2 supersedes IcePatch, which was a previous version of this service.

IcePatch2 is an efficient file patching service that is easy to configure and use. It includes the following components:

- the IcePatch server (`icepatch2server`)
- a text-based IcePatch client (`icepatch2client`)
- a text-based tool to compress files and calculate checksums (`icepatch2calc`)
- a Slice API and C++ convenience library for developing custom IcePatch2 clients

As with all Ice services, IcePatch2 can be configured to use Ice facilities such as Glacier2 for firewall support and IceSSL for secure communication.

IcePatch2 is conceptually quite simple. The server is given responsibility for a file system directory (the *data directory*) containing the files and subdirectories that are to be distributed to IcePatch2 clients. You use `icepatch2calc` to compress these files and to generate an index containing a checksum for each file. The server transmits the compressed files to the client, which recreates the data directory and its contents on the client side, patching any files that have changed since the previous run.

IcePatch2 is efficient: transfer rates for files are comparable to what you would get using `ftp`.

IcePatch2 addresses a requirement common to both development and deployment scenarios: the safe, secure, and efficient replication of a directory tree. The IcePatch2 server is easy to configure and efficient. For simple uses, IcePatch2 provides a client that can be used to patch directory hierarchies from the command line. With the C++ utility library, you can also create custom patch clients if you require better integration of the client with your application.

Topics

- Using icepatch2calc
- Running the IcePatch2 Server
- Running the IcePatch2 Client
- IcePatch2 Object Identities
- IcePatch2 Client Utility Library

Copyright 2016, ZeroC, Inc.

Using icepatch2calc

This page describes how to prepare a file set using `icepatch2calc`.

On this page:

- Preparing a File Set using icepatch2calc
- icepatch2calc Command Line Options

Preparing a File Set using `icepatch2calc`

Suppose we have the directories and files shown below:

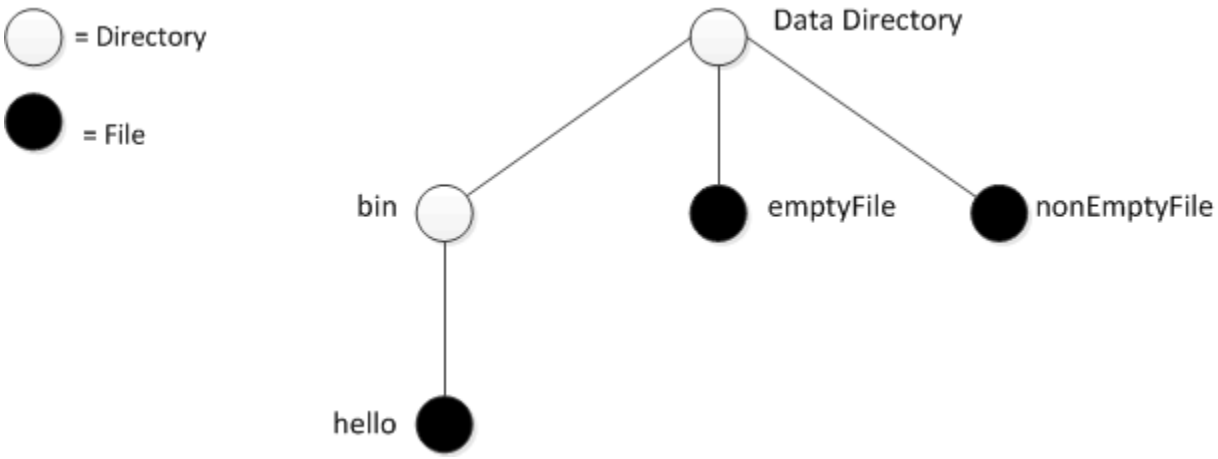

An example data directory.

Assume that the file named `emptyFile` is empty (contains zero bytes) and that the remaining files contain data.

To prepare this directory for the transmission by the server, you must first run `icepatch2calc`. (The command shown assumes that the data directory is the current directory.)

```
$ icepatch2calc .
```

After running this command, the contents of the data directory look as follows:

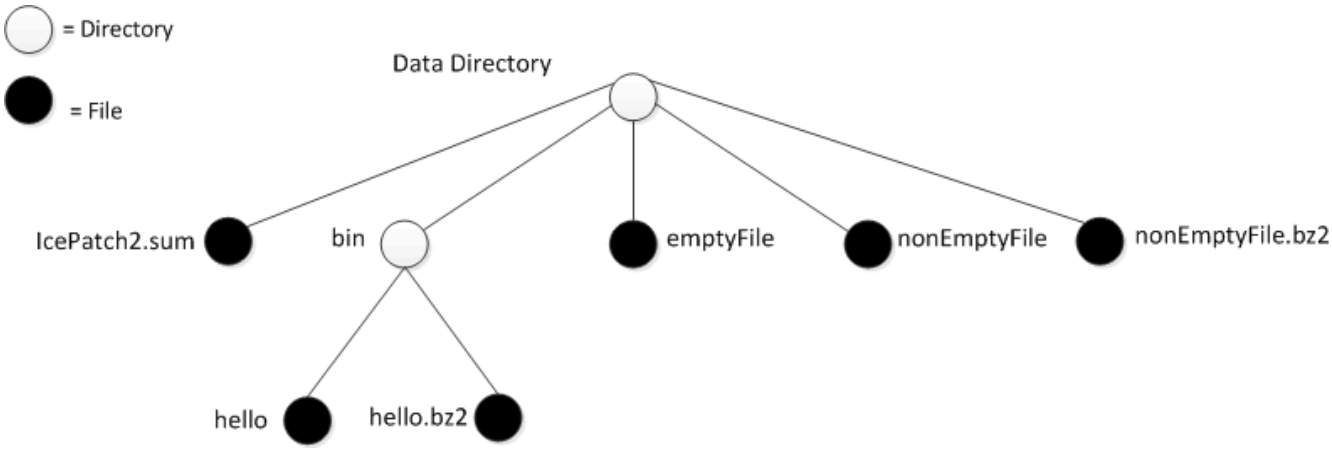

Contents of the data directory after running `icepatch2calc`.

Note that `icepatch2calc` compresses the files in the data directory (except for `emptyFile`, which is not compressed). Also note that `icepatch2calc` creates an additional file, `IcePatch2.sum` in the data directory. The contents of this file are as follows:

Copyright 2016, ZeroC, Inc.

```
. 3a52ce780950d4d969792a2559cd519d7ee8c727 -1
./bin bd362140a3074eb3edb5e4657561e029092c3d91 -1
./bin/hello 77b11db586a1f20aab8553284241bb3cd532b3d5 70
./emptyFile 082c37fc2641db68d195df83844168f8a464eada 0
./nonEmptyFile aec7301c408e6ce184ae5a34e0ea46e0f0563746 72
```

Each line in the checksum file contains the name of the *uncompressed* file or directory (relative to the data directory), the checksum of the *uncompressed* file, and a byte count. For directories, the count is `-1`; for uncompressed files, the count is `0`; for compressed files, the count is the number of bytes in the *compressed* file. The lines in the file are sorted alphabetically by their pathname.

If you add files or delete files from the data directory or make changes to existing files, you must stop the server, run `icepatch2calc` again to update the `IcePatch2.sum` checksum file, and restart the server.

`icepatch2calc` Command Line Options

`icepatch2calc` has the following syntax:

```
icepatch2calc [options] data_dir [file...]
```

Normally, you will run `icepatch2calc` by simply specifying a data directory, in which case the program traverses the data directory, compresses all files, and creates an entry in the checksum file for each file and directory.

You can also nominate specific files or directories on the command line. In this case, `icepatch2calc` only compresses and calculates checksums for the specified files and directories. This is useful if you have a very large file tree and want to refresh the checksum entries for only a few selected files or directories that you have updated. (In this case, the program does not traverse the entire data directory and, therefore, will also not detect any updated, added, or deleted files, except in any of the specified directories.) Any file or directory names you specify on the command line must either be pathnames relative to the data directory or, if you use absolute pathnames, those pathnames must have the data directory as a prefix.

The command supports the following options:

- `-h, --help`
  Displays a help message.
- `-v, --version`
  Displays the version number.
- `-z, --compress`
  Normally, `icepatch2calc` scans the data directory and compresses a file only if no compressed version exists, or if the compressed version of a file has a modification time that predates that of the uncompressed version. If you specify `-z`, the tool re-scans and recompresses the entire data directory, regardless of the time stamps on files. This option is useful if you suspect that time stamps in the data directory may be incorrect.
- `-Z, --no-compress`
  This option allows you to create a client-side checksum file. Do not use this option when creating the checksum file for the server — the option is for creating a client-side `IcePatch2.sum` file for updates of software on distribution media.
- `-i, --case-insensitive`
  This option disallows file names that differ only in case. (An error message will be printed if `icepatch2calc` encounters any files that differ in case only.) This is particularly useful for Unix servers with Windows clients, since Windows folds the case of file names, and therefore such files would override each other on the Windows client.
- `-V, --verbose`
  This option prints a progress message for each file that is compressed and for each checksum that is computed.

See Also

- Running the IcePatch2 Client

Copyright 2016, ZeroC, Inc.

Running the IcePatch2 Server

This page describes how to run the IcePatch2 server.

On this page:

- Starting icepatch2server
- icepatch2server Command Line Options

Starting `icepatch2server`

Once you have run `icepatch2calc` on the data directory, you can start the `icepatch2server`:

```
$ icepatch2server .
```

The server expects the data directory as its single command-line argument. If you omit to specify the data directory, the server uses the setting of the `IcePatch2.Directory` property to determine the data directory.

You must also specify the endpoints at which the server listens for client requests, by setting the `IcePatch2.Endpoints` property.

`icepatch2server` Command Line Options

Regardless of whether you run the server under Windows or a Unix-like operating system, it provides the following options:

- `-h, --help`
  Displays a help message.
- `-v, --version`
  Displays a version number.

Additional command line options are supported, including those that allow the server to run as a Windows service or Unix daemon.

See Also

- The Server-Side main Function in C++

Copyright 2016, ZeroC, Inc.

Running the IcePatch2 Client

This page describes how to un the IcePatch2 client.

On this page:

- Patching with icepatch2client
- Using icepatch2client for Partial Updates
- Preventing Deletion of Local Files
- Patching Software Installed from Media
- Setting Transfer Size for icepatch2client
- icepatch2client Command Line Options

Patching with `icepatch2client`

Once the `icepatch2server` is running, you can use `icepatch2client` to get a copy of the data directory that is maintained by the server. For example:

```
$ icepatch2client --IcePatch2Client.Proxy="IcePatch2Server/server:tcp -h
somehost.com -p 10000" .
```

The client expects the data directory as its single command-line argument. You must provide a proxy to the IcePatch2 server with the `IcePatch2Client.Proxy` property, as shown above.

If you have not run the client previously, it asks you whether you want to do a thorough patch. You must reply "yes" at this point (or run the client with the `-t` option to perform a thorough update). The client then executes the following steps:

1. It traverses the local data directory and creates a local `IcePatch2.sum` checksum file.
2. It obtains the relevant list of checksums from the server and compares it to the list of checksums it has built locally:
    - The client deletes each file that appears in the local checksum file but not in the server's file.
    - The client retrieves every file that appears in the server's checksum file, but not in the local checksum file.
    - The client patches every file that, locally, has a checksum that differs from the corresponding checksum on the server side.

When the client finishes, the contents of the data directory on the client side exactly match the contents of the data directory on the server side. However, only the uncompressed files are created on the client side — the server stores the compressed version of the files simply to avoid redundantly compressing a file every time it is retrieved by a client.

On the initial patch, any files that exist in the client's data directory are deleted or, if they have the same name as a file on the server, will be overwritten with the corresponding file as it exists on the server.

Using `icepatch2client` for Partial Updates

Once you have run the client, the client-side data directory contains an `IcePatch2.sum` file that reflects the contents of the data directory. If you run `icepatch2client` a second time, the program uses the contents of the local checksum file: for each entry in the local checksum file, the client compares the local checksum with the server-side checksum for the same file; if the checksums differ, the client updates the corresponding file. In addition, the client deletes any files that appear in the client's checksum file but not in the server's checksum file, and it fetches any files that appear in the server's checksum file but are missing from the client's checksum file.

If you edit a client-side file and change its contents, `icepatch2client` does *not* realize that this has happened and therefore will not patch the file to be in sync with the version on the server again. This is because the client does not automatically recompute the checksum for a file to see whether the stored checksum in `IcePatch2.sum` still agrees with the actual checksum for the current file contents.

Similarly, if you create an arbitrary file on the client side, but that file is not mentioned in either the client's or the server's checksum file, that file will simply be left alone. In other words, a normal patch operates on the differences between the client's and server's checksum files, not on any differences that could be detected by examining the contents of the file system.

If you have locally created files that have nothing to do with the distribution or if you have locally modified some files and want to make sure that those modified files are updated to reflect the contents of the same files on the server side, you must run a thorough patch with the `-t` option. This forces the client to traverse the local data directory and recompute the checksum for each file, and then compare these

Copyright 2016, ZeroC, Inc.

checksums against the server-side ones. As a result, if you edit a file locally so it differs from the server-side version, `-t` forces that file to be updated. Similarly, if you have added a file of your own on the client side that does not have a counterpart on the server side, that file will be deleted by a thorough patch.

Preventing Deletion of Local Files

By default, a normal patch deletes any files that appear in the client's checksum file but that are absent in the server's checksum file. Similarly, by default, a thorough patch deletes all files in the local data directory that do not appear in the server's checksum file. If you do not want this behavior, you can set the `IcePatch2Client.Remove` property to 0 (the default value is 1). This prevents deletion of files and directories that exist only on the client side, whether the patch is a normal patch or a thorough patch.

Patching Software Installed from Media

Suppose you distribute your application on a DVD that clients use to install the software. The DVD might be out of date so, after installation, the install script needs to perform a patch to update the application to the latest version. The script can perform a thorough patch to do this but, for large file sets, this is expensive because the client has to recompute the checksum for every file in the distribution.

To avoid this cost, you can place all the files for the distribution into a directory on the server and run `icepatch2calc -Z` on that directory. With the `-Z` option, `icepatch2calc` creates a checksum file with the correct checksums, but with a file size of 0 for each file, that is, the `-Z` option omits compressing the files (and the considerable cost associated with that). Once you have created the new `IcePatch2.sum` file in this way, you can include it on the DVD and install it on the client along with all the other files.

This guarantees that the checksum file on the client is in agreement with the actual files that were just installed and, therefore, it is sufficient for the install script to do a normal patch to update the distribution and so avoid the cost of recomputing the checksum for every file.

Setting Transfer Size for `icepatch2client`

You can set the `IcePatch2Client.ChunkSize` property to control the number of bytes that the client fetches per request. The default value is 100 kilobytes.

`icepatch2client` Command Line Options

The client supports the following options:

- `-h, --help`
  Displays a help message.
- `-v, --version`
  Displays a version number.
- `-t, --thorough`
  Do a thorough patch, recomputing all checksums.

See Also

- Running the IcePatch2 Server
- IcePatch2Client.*

Copyright 2016, ZeroC, Inc.

IcePatch2 Object Identities

An IcePatch2 service hosts one well-known object, which implements the `IcePatch2::FileServer` interface and has the default identity `IcePatch2/server`. If an application requires the use of multiple IcePatch2 services, it is a good idea to assign unique identities to the well-known objects by configuring the services with different values for the `IcePatch2.InstanceName` property, as shown in the following example:

```
$ icepatch2server --IcePatch2.InstanceName=PublicFiles ...
```

This property changes the category of the server object's identity, which becomes `PublicFiles/server`.

See Also

- Object Identity
- IcePatch2.*

Copyright 2016, ZeroC, Inc.

IcePatch2 Client Utility Library

IcePatch2 includes a pair of C++ classes that simplify the task of writing your own patch client, along with a Microsoft Foundation Classes (MFC) example that shows how to use these classes. You can find the MFC example in the subdirectory `demo/IcePatch2/MFC` of your Ice distribution.

The remainder of this section discusses the classes. To incorporate them into a custom patch client, your program must include the header file `IcePatch2/ClientUtil.h` and link with the `IcePatch2` library.

On this page:

- Performing a Patch
  - Constructing a Patcher
  - Executing the Patch
- Monitoring Patch Progress

Performing a Patch

The `Patcher` class encapsulates all of the patching logic required by a client:

C++

```
namespace IcePatch2 {
class Patcher : ... {
public:

    Patcher(const Ice::CommunicatorPtr& communicator,
            const PatcherFeedbackPtr& feedback);

    Patcher(const FileServerPrx& server,
            const PatcherFeedbackPtr& feedback,
            const std::string& dataDir, bool thorough,
            Ice::Int chunkSize, Ice::Int remove);

    bool prepare();

    bool patch(const std::string& dir);

    void finish();
};
typedef IceUtil::Handle<Patcher> PatcherPtr;
}
```

Constructing a Patcher

The constructors provide two ways of configuring a `Patcher` instance. The first form obtains the following IcePatch2Client configuration properties from the supplied communicator:

- `IcePatch2Client.Proxy`
- `IcePatch2Client.Directory`
- `IcePatch2Client.Thorough`
- `IcePatch2Client.ChunkSize`
- `IcePatch2Client.Remove`

The second constructor accepts arguments that correspond to each of these properties.

Both constructors also accept a `PatcherFeedback` object, which allows the client to monitor the progress of the patch.

Copyright 2016, ZeroC, Inc.

Executing the Patch

`Patcher` provides three methods that reflect the three stages of a patch:

- `bool prepare()`
  The first stage of a patch includes reading the contents of the checksum file (if present), retrieving the file information from the server, and examining the local data directory to compose the list of files that require updates. The `PatcherFeedback` object is notified incrementally about each local task and has the option of aborting the patch at any time. This method returns true if patch preparation completed successfully, or false if the `PatcherFeedback` object aborted the patch. If an error occurs, `prepare` raises an exception in the form of a `std::string` containing a description of the problem.

- `bool patch(const std::string& dir)`
  The second stage of a patch updates the files in the local data directory. If the `dir` argument is an empty string or `"."`, `patch` updates the entire data directory. Otherwise, `patch` updates only those files whose path names begin with the path in `dir`. For each file requiring an update, `Patcher` downloads its compressed data from the server and writes it to the local data directory. The `PatcherFeedback` object is notified about the progress of each file and, as in the preparation stage, may abort the patch if necessary. This method returns true if patching completed successfully, or false if the `PatcherFeedback` object aborted the patch. If an error occurs, `patch` raises an exception in the form of a `std::string` containing a description of the problem.

- `void finish()`
  The final stage of a patch writes a new checksum file to the local data directory. If an error occurs, `finish` raises an exception in the form of a `std::string` containing a description of the problem.

The code below demonstrates a simple patch client:

C++

```
#include <IcePatch2/ClientUtil.h>
...
Ice::CommunicatorPtr communicator = ...;
IcePatch2::PatcherFeedbackPtr feedback = new MyPatcherFeedbackI;
IcePatch2::PatcherPtr patcher =
    new IcePatch2::Patcher(communicator, feedback);

try {
    bool aborted = !patcher->prepare();
    if(!aborted)
        aborted = !patcher->patch("");
    if(!aborted)
        patcher->finish();
    if(aborted)
        cerr << "Patch aborted" << endl;
} catch(const string& reason) {
    cerr << "Patch error: " << reason << endl;
}
```

For a more sophisticated example, see `demo/IcePatch2/MFC` in your Ice distribution.

Monitoring Patch Progress

The class `PatcherFeedback` is an abstract base class that allows you to monitor the progress of a `Patcher` object. The class declaration is shown below:

Copyright 2016, ZeroC, Inc.

C++

```
namespace IcePatch2 {
class PatcherFeedback : ... {
public:

    virtual bool noFileSummary(const std::string& reason) = 0;

    virtual bool checksumStart() = 0;
    virtual bool checksumProgress(const std::string& path) = 0;
    virtual bool checksumEnd() = 0;

    virtual bool fileListStart() = 0;
    virtual bool fileListProgress(Ice::Int percent) = 0;
    virtual bool fileListEnd() = 0;

    virtual bool patchStart(
        const std::string& path, Ice::Long size,
        Ice::Long updated, Ice::Long total) = 0;
    virtual bool patchProgress(
        Ice::Long pos, Ice::Long size,
        Ice::Long updated, Ice::Long total) = 0;
    virtual bool patchEnd() = 0;
};
typedef IceUtil::Handle<PatcherFeedback> PatcherFeedbackPtr;
}
```

Each of these methods returns a boolean value:

- true allows `Patcher` to continue
- false directs `Patcher` to abort the patch.

The methods are described below.

- `bool noFileSummary(const std::string& reason)`
  Invoked when the local checksum file cannot be found. Returning true initiates a thorough patch, while returning false causes `Patcher::prepare` to return false.

- `bool checksumStart()`
  `bool checksumProgress(const std::string& path)`
  `bool checksumEnd()`
  Invoked by `Patcher::prepare` during a thorough patch. The `checksumProgress` method is invoked as each file's checksum is being computed.

- `bool fileListStart()`
  `bool fileListProgress(Ice::Int percent)`
  `bool fileListEnd()`
  Invoked by `Patcher::prepare` when collecting the list of files to be updated. The `percent` argument to `fileListProgress` indicates how much of the collection process has completed so far.

- `bool patchStart(const std::string& path, Ice::Long size, Ice::Long updated, Ice::Long total)`
  `bool patchProgress(Ice::Long pos, Ice::Long size, Ice::Long updated, Ice::Long total)`
  `bool patchEnd()`
  For each file that requires updating, `Patcher::patch` invokes `patchStart` to indicate the beginning of the patch, `patchProgress` one or more times as chunks of the file are downloaded and written, and finally `patchEnd` to signal the completion of the file's

Copyright 2016, ZeroC, Inc.

patch. The `path` argument supplies the path name of the file, and `size` provides the file's compressed size. The `pos` argument denotes the number of bytes written so far, while `updated` and `total` represent the cumulative number of bytes updated so far and the total number of bytes to be updated, respectively, of the entire patch operation.

See Also

- IcePatch2Client.*

Copyright 2016, ZeroC, Inc.

IceStorm

IceStorm is an efficient publish/subscribe service for Ice applications. Applications often need to disseminate information to multiple recipients. For example, suppose we are developing a weather monitoring application in which we collect measurements such as wind speed and temperature from a meteorological tower and periodically distribute them to weather monitoring stations. We initially consider using the architecture shown below:

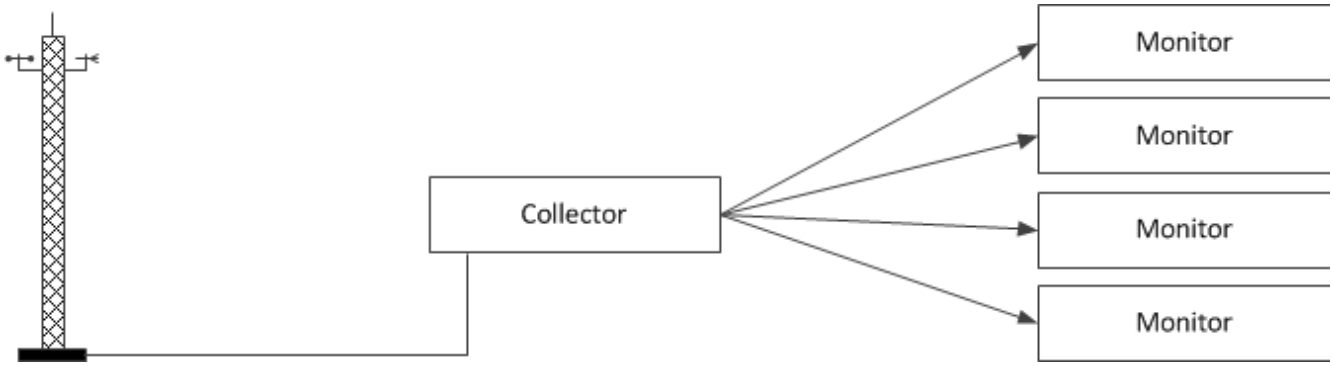

Initial design for a weather monitoring application.

However, the primary disadvantage of this architecture is that it tightly couples the collector to its monitors, needlessly complicating the collector implementation by requiring it to manage the details of monitor registration, measurement delivery, and error recovery. We can rid ourselves of these mundane duties by incorporating IceStorm into our architecture, as shown below:

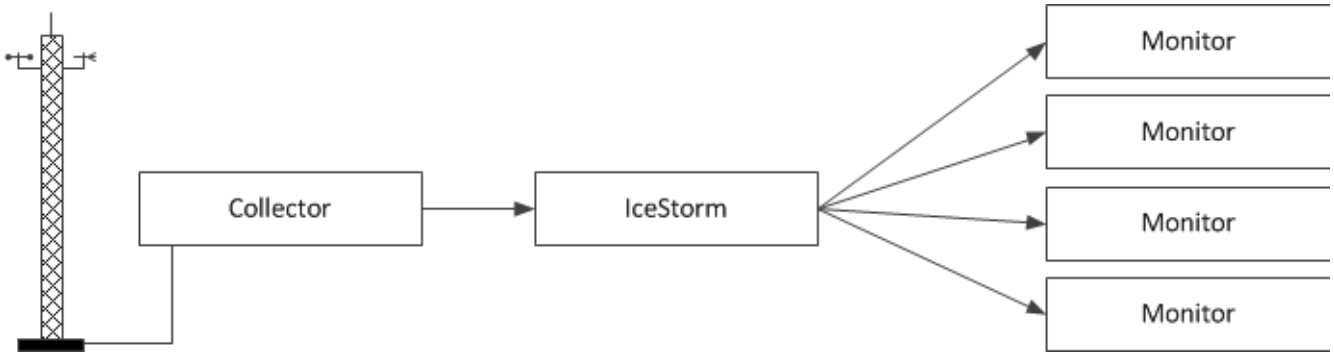

A weather monitoring application using IceStorm.

IceStorm simplifies the collector implementation significantly by decoupling it from the monitors. As a publish/subscribe service, IceStorm acts as a mediator between the collector (the publisher) and the monitors (the subscribers), and offers several advantages:

- When the collector is ready to distribute a new set of measurements, it makes a single request to the IceStorm server. The IceStorm server takes responsibility for delivering the request to the monitors, including handling any exceptions caused by ill-behaved or missing subscribers. The collector no longer needs to be aware of its monitors, or whether it even has any monitors at that moment.
- Similarly, monitors interact with the IceStorm server to perform tasks such as subscribing and unsubscribing, thereby allowing the collector to focus on its application-specific responsibilities and not on administrative trivia.
- The collector and monitor applications require very few changes to incorporate IceStorm.

Topics

- IceStorm Concepts
- IceStorm Interfaces
- Using IceStorm
- Highly Available IceStorm
- IceStorm Administration
- Topic Federation
- IceStorm Quality of Service
- IceStorm Delivery Modes
- Configuring IceStorm
- IceStorm Metrics
- IceStorm Database Utility

Copyright 2016, ZeroC, Inc.

IceStorm Concepts

This section discusses several concepts that are important for understanding IceStorm's capabilities.

- **Message**
  An IceStorm *message* is strongly typed and is represented by an invocation of a Slice operation: the operation name identifies the type of the message, and the operation parameters define the message contents. A message is published by invoking the operation on an IceStorm proxy in the normal fashion. Similarly, subscribers receive the message as a regular servant upcall. As a result, IceStorm uses the "push" model for message delivery; polling is not supported.

- **IceStorm Topics**
  An application indicates its interest in receiving messages by subscribing to a *topic*. An IceStorm server supports any number of topics, which are created dynamically and distinguished by unique names. Each topic can have multiple publishers and subscribers.

  A topic is essentially equivalent to an application-defined Slice interface: the operations of the interface define the types of messages supported by the topic. A publisher uses a proxy for the topic interface to send its messages, and a subscriber implements the topic interface (or an interface derived from the topic interface) in order to receive the messages. This is no different than if the publisher and subscriber were communicating directly in the traditional client-server style; the interface represents the contract between the client (the publisher) and the server (the subscriber), except IceStorm transparently forwards each message to multiple recipients.

  IceStorm does not verify that publishers and subscribers are using compatible interfaces, therefore applications must ensure that topics are used correctly.

- **Unidirectional Messages**
  IceStorm messages are unidirectional, that is, they must have `void` return type, cannot have out-parameters, and cannot raise user exceptions. It follows that a publisher cannot receive replies from its subscribers. Any of the Ice transports (TCP, SSL, and UDP) can be used to publish and receive messages.

- **Federation**
  IceStorm supports the formation of topic graphs, also known as federation. A topic graph is formed by creating links between topics, where a *link* is a unidirectional association from one topic to another. Each link has a *cost* that may restrict message delivery on that link. A message published on a topic is also published on all of the topic's links for which the message cost does not exceed the link cost.

  Once a message has been published on a link, the receiving topic publishes the message to its subscribers, but does not publish it on any of its links. In other words, IceStorm messages propagate at most one hop from the originating topic in a federation.

  The following figure presents an example of topic federation. Topic T_1 has links to T_2 and T_3, as indicated by the arrows. The subscribers S_1 and S_2 receive all messages published on T_2, as well as those published on T_1. Subscriber S_3 receives messages only from T_1, and S_4 receives messages from both T_3 and T_1.

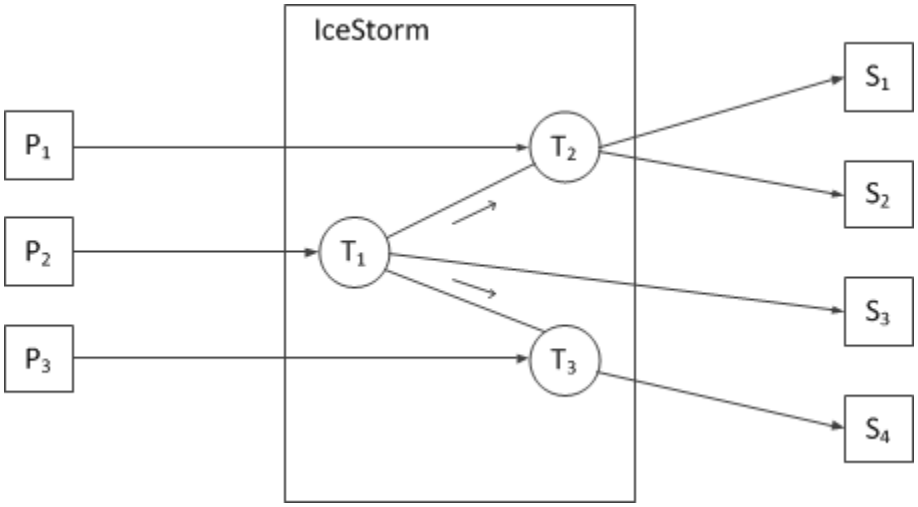

Topic federation.

  IceStorm makes no attempt to prevent a subscriber from receiving duplicate messages. For example, if a subscriber is subscribed to both T_2 and T_3, then it would receive two requests for each message published on T_1.

- **Quality of Service**

Copyright 2016, ZeroC, Inc.

IceStorm allows each subscriber to specify its own *quality of service* (QoS) parameters that affect the delivery of its messages. Quality of service parameters are represented as a dictionary of name-value pairs.

- **Replication**
  IceStorm supports replication to provide higher availability for publishers and subscribers.
- **Persistent Mode**
  IceStorm's default behavior maintains information about topics, links, and subscribers in a database. However, a message sent via IceStorm is not stored persistently, but rather is discarded as soon as it is delivered to the topic's current set of subscribers. If an error occurs during delivery to a subscriber, IceStorm does not queue messages for that subscriber.
- **Transient Mode**
  IceStorm can optionally run in a fully transient mode in which no database is required. Replication is not supported in this mode.
- **Subscriber Errors**

  IceStorm automatically removes a subscription from a topic if a subscriber failure occurs while attempting to deliver a message. For example, IceStorm may be unable to establish a connection to the subscriber using the proxy that the subscriber provided, meaning the subscriber is not currently active at the proxy's endpoints, or those endpoints are inaccessible to IceStorm. Another common failure scenario is a subscriber that allows an exception to propagate back to IceStorm. This is important if you make changes to a Slice data type or operation signature: if you do, you must ensure that both publishers and subscribers use the same Slice definitions; if you do not, the subscriber is likely to encounter marshaling errors when receiving an event from IceStorm with a mismatched Slice definition. If the subscriber allows this error to propagate back to IceStorm, its subscription will be canceled.

  Use the `retryCount` quality of service parameter to configure IceStorm's behavior in error situations.

See Also

- Oneway Invocations
- Topic Federation
- IceStorm Quality of Service
- Configuring IceStorm
- Freeze

Copyright 2016, ZeroC, Inc.

IceStorm Interfaces

This page provides a brief introduction to the Slice interfaces comprising the IceStorm service. See the online Slice API Reference XREF for the Slice documentation.

On this page:

- The TopicManager Interface
- The Topic Interface

The `TopicManager` Interface

The `TopicManager` is a singleton object that acts as a factory and repository of `Topic` objects. Its interface and related types are shown below:

Slice

```
module IceStorm {
dictionary<string, Topic*> TopicDict;

exception TopicExists {
    string name;
};

exception NoSuchTopic {
    string name;
};

interface TopicManager {
    Topic* create(string name) throws TopicExists;
    idempotent Topic* retrieve(string name) throws NoSuchTopic;
    idempotent TopicDict retrieveAll();
    idempotent Ice::SliceChecksumDict getSliceChecksums();
};
};
```

The `create` operation is used to create a new topic, which must have a unique name. The `retrieve` operation allows a client to obtain a proxy for an existing topic, and `retrieveAll` supplies a dictionary of all existing topics. The `getSliceChecksums` operation returns Slice checksums for the IceStorm definitions.

The `Topic` Interface

The `Topic` interface represents a topic and provides several administrative operations for configuring links and managing subscribers.

Copyright 2016, ZeroC, Inc.

Slice

```
module IceStorm {
struct LinkInfo {
    Topic* theTopic;
    string name;
    int cost;
};
sequence<LinkInfo> LinkInfoSeq;

dictionary<string, string> QoS;

exception LinkExists {
    string name;
};

exception NoSuchLink {
    string name;
};

exception AlreadySubscribed {};

exception BadQoS {
    string reason;
};

interface Topic {
    idempotent string getName();
    idempotent Object* getPublisher();
    idempotent Object* getNonReplicatedPublisher();
    Object* subscribeAndGetPublisher(QoS theQoS, Object* subscriber)
        throws AlreadySubscribed, BadQoS;
    idempotent void unsubscribe(Object* subscriber);
    idempotent void link(Topic* linkTo, int cost)
        throws LinkExists;
    idempotent void unlink(Topic* linkTo) throws NoSuchLink;
    idempotent LinkInfoSeq getLinkInfoSeq();
    void destroy();
};
};
```

The `getName` operation returns the name assigned to the topic, while the `getPublisher` and `getNonReplicatedPublisher` operations return proxies for the topic's publisher object.

The `subscribeAndGetPublisher` operation adds a subscriber's proxy to the topic; if another subscriber proxy already exists with the same object identity, the operation throws `AlreadySubscribed`. The operation returns a proxy for a subscriber-specific publisher object.

The `unsubscribe` operation removes the subscriber from the topic.

A link to another topic is created using the `link` operation; if a link already exists to the given topic, the `LinkExists` exception is raised. Links are destroyed using the `unlink` operation.

Finally, the `destroy` operation permanently destroys the topic.

See Also

Copyright 2016, ZeroC, Inc.

- Slice Checksums
- Using an IceStorm Publisher Object
- Publishing to a Specific Subscriber
- Topic Federation

Copyright 2016, ZeroC, Inc.

Using IceStorm

Now we'll expand on the earlier weather monitoring example, demonstrating how to create, subscribe to and publish messages on a topic. We use the following Slice definitions in our example:

Slice

```
struct Measurement {
    string tower; // tower id
    float windSpeed; // knots
    short windDirection; // degrees
    float temperature; // degrees Celsius
};

interface Monitor {
    void report(Measurement m);
};
```

`Monitor` is our topic interface. For the sake of simplicity, it defines just one operation, `report`, taking a `Measurement` struct as its only parameter.

Topics

- Implementing an IceStorm Publisher
- Using an IceStorm Publisher Object
- Implementing an IceStorm Subscriber
- Publishing to a Specific Subscriber

See Also

- IceStorm

Copyright 2016, ZeroC, Inc.

Implementing an IceStorm Publisher

The implementation of our weather measurement collector application can be summarized easily:

1. Obtain a proxy for the `TopicManager`. This is the primary IceStorm object, used by both publishers and subscribers.
2. Obtain a proxy for the `Weather` topic, either by creating the topic if it does not exist, or retrieving the proxy for the existing topic.
3. Obtain a proxy for the `Weather` topic's "publisher object." This proxy is provided for the purpose of publishing messages, and therefore is narrowed to the topic interface (`Monitor`).
4. Collect and report measurements.

We present collector implementations in C++ and Java below.

On this page:

- Publisher Example in C++
- Publisher Example in Java

Publisher Example in C++

As usual, our C++ example begins by including the necessary header files. The interesting ones are `IceStorm/IceStorm.h`, which is generated from the IceStorm Slice definitions, and `Monitor.h`, containing the generated code for our monitor definitions shown above.

C++

```
#include <Ice/Ice.h>
#include <IceStorm/IceStorm.h>
#include <Monitor.h>

int main(int argc, char* argv[])
{
    ...
    Ice::ObjectPrx obj = communicator->stringToProxy("IceStorm/TopicManager:tcp -p
9999");
    IceStorm::TopicManagerPrx topicManager =
IceStorm::TopicManagerPrx::checkedCast(obj);
    IceStorm::TopicPrx topic;
    while (!topic) {
        try {
            topic = topicManager->retrieve("Weather");
        } catch (const IceStorm::NoSuchTopic&) {
            try {
                topic = topicManager->create("Weather");
            } catch (const IceStorm::TopicExists&) {
                // Another client created the topic.
            }
        }
    }

    Ice::ObjectPrx pub = topic->getPublisher()->ice_oneway();
    MonitorPrx monitor = MonitorPrx::uncheckedCast(pub);
    while (true) {
        Measurement m = getMeasurement();
        monitor->report(m);
    }
    ...
}
```


Copyright 2016, ZeroC, Inc.

Note that this example assumes that IceStorm uses the instance name `IceStorm`. The actual instance name may differ, and you need to use it as the category when calling `stringToProxy`.

After obtaining a proxy for the topic manager, the collector attempts to retrieve the topic. If the topic does not exist yet, the collector receives a `NoSuchTopic` exception and then creates the topic:

C++

```
IceStorm::TopicPrx topic;
while(!topic) {
    try {
        topic = topicManager->retrieve("Weather");
    } catch (const IceStorm::NoSuchTopic&) {
        try {
            topic = topicManager->create("Weather");
        } catch (const IceStorm::TopicExists&) {
            // Another client created the topic.
        }
    }
}
```

The next step is obtaining a proxy for the publisher object, which the collector narrows to the `Monitor` interface. (We create a oneway proxy for the publisher purely for efficiency reasons.)

C++

```
Ice::ObjectPrx pub = topic->getPublisher()->ice_oneway();
MonitorPrx monitor = MonitorPrx::uncheckedCast(pub);
```

Finally, the collector enters its main loop, collecting measurements and publishing them via the IceStorm publisher object:

C++

```
while (true) {
    Measurement m = getMeasurement();
    monitor->report(m);
}
```

Publisher Example in Java

The equivalent Java version is shown below:

Copyright 2016, ZeroC, Inc.

Java

```
public static void main(String[] args)
{
    ...
    Ice.ObjectPrx obj = communicator.stringToProxy("IceStorm/TopicManager:tcp -p 99
99");
    IceStorm.TopicManagerPrx topicManager =
IceStorm.TopicManagerPrxHelper.checkedCast(obj);
    IceStorm.TopicPrx topic = null;
    while (topic == null) {
        try {
            topic = topicManager.retrieve("Weather");
        } catch (IceStorm.NoSuchTopic ex) {
            try {
                topic = topicManager.create("Weather");
            } catch (IceStorm.TopicExists ex) {
                // Another client created the topic.
            }
        }
    }

    Ice.ObjectPrx pub = topic.getPublisher().ice_oneway();
    MonitorPrx monitor = MonitorPrxHelper.uncheckedCast(pub);
    while (true) {
        Measurement m = getMeasurement();
        monitor.report(m);
    }
    ...
}
```

Note that this example assumes that IceStorm uses the instance name `IceStorm`. The actual instance name may differ, and you need to use it as the category when calling `stringToProxy`.

After obtaining a proxy for the topic manager, the collector attempts to retrieve the topic. If the topic does not exist yet, the collector receives a `NoSuchTopic` exception and then creates the topic:

Copyright 2016, ZeroC, Inc.

Java

```
IceStorm.TopicPrx topic = null;
while (topic == null) {
    try {
        topic = topicManager.retrieve("Weather");
    } catch (IceStorm.NoSuchTopic ex) {
        try {
            topic = topicManager.create("Weather");
        } catch (IceStorm.TopicExists ex) {
            // Another client created the topic.
        }
    }
}
```

The next step is obtaining a proxy for the publisher object, which the collector narrows to the `Monitor` interface:

Java

```
Ice.ObjectPrx pub = topic.getPublisher().ice_oneway();
MonitorPrx monitor = MonitorPrxHelper.uncheckedCast(pub);
```

Finally, the collector enters its main loop, collecting measurements and publishing them via the IceStorm publisher object:

Java

```
while (true) {
    Measurement m = getMeasurement();
    monitor.report(m);
}
```

See Also

- Configuring IceStorm

Copyright 2016, ZeroC, Inc.

Using an IceStorm Publisher Object

Each topic creates a publisher object for the express purpose of publishing messages. It is a special object in that it implements an Ice interface that allows the object to receive and forward requests (i.e., IceStorm messages) without requiring knowledge of the operation types.

On this page:

- Type Safety Considerations for the Publisher Object
- Publish using Oneway or Twoway Invocations?
- Selecting a Transport for the Publisher Object
- Using Request Contexts with the Publisher Object

Type Safety Considerations for the Publisher Object

From the publisher's perspective, the publisher object appears to be an application-specific type. In reality, the publisher object can forward requests for any type, and that introduces a degree of risk: a misbehaving publisher can use `uncheckedCast` to narrow the publisher object to any type and invoke any operation; the publisher object unknowingly forwards those requests to the subscribers.

If a publisher sends a request using an incorrect type, the Ice run time in a subscriber typically responds by raising `OperationNotExistException`. However, since the subscriber receives its messages as oneway invocations, no response can be sent to the publisher object to indicate this failure, and therefore neither the publisher nor the subscriber is aware of the type-mismatch problem. In short, IceStorm places the burden on the developer to ensure that publishers and subscribers are using it correctly.

Publish using Oneway or Twoway Invocations?

IceStorm messages are unidirectional, but publishers may use either oneway or twoway invocations when sending messages to the publisher object. Each invocation style has advantages and disadvantages that you should consider when deciding which one to use. The differences between the invocation styles affect a publisher in four ways:

- Efficiency
  Oneway invocations have the advantage in efficiency because the Ice run time in the publisher does not await a reply to each message (and, of course, no reply is sent by IceStorm on the wire).
- Ordering
  The use of oneway invocations by a publisher may affect the order in which subscribers receive messages. If ordering is important, use twoway invocations with a reliability QoS of `ordered`, or use a single thread in the subscriber.
- Reliability
  Oneway invocations can be lost under certain circumstances, even when they are sent over a reliable transport such as TCP. If the loss of messages is unacceptable, or you are unable to address the potential causes of lost oneway messages, then twoway invocations are recommended.
- Delays
  A publisher may experience network-related delays when sending messages to IceStorm if subscribers are slow in processing messages. Twoway invocations are more susceptible to these delays than oneway invocations.

Selecting a Transport for the Publisher Object

Each publisher can select its own transport for message delivery, therefore the transport used by a publisher to communicate with IceStorm has no effect on how IceStorm delivers messages to its subscribers.

For example, a publisher can use a UDP transport if the possibility of lost messages is acceptable (and if IceStorm provides a UDP endpoint to publishers). However, the TCP or SSL transports are generally recommended for IceStorm's publisher endpoint in order to ensure that published messages are delivered reliably to IceStorm, even if they may not be delivered reliably to some subscribers.

Using Request Contexts with the Publisher Object

A request context is an optional argument of all remote invocations. If a publisher supplies a request context when publishing a message, IceStorm will forward it intact to subscribers.

Services such as Glacier2 employ request contexts to provide applications with more control over the service's behavior. For example, if a publisher knows that IceStorm is delivering messages to subscribers via a Glacier2 router, the publisher can influence Glacier2's behavior by including a request context, as shown in the following C++ example:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx pub = topic->getPublisher();
Ice::Context ctx;
ctx["_fwd"] = "Oz";
MonitorPrx monitor = MonitorPrx::uncheckedCast(pub->ice_context(ctx));
```

The `_fwd` context key, when encountered by Glacier2, causes the router to forward the request using compressed batch oneway messages. The `ice_context` proxy method is used to obtain a proxy that includes the Glacier2 request context in every invocation, eliminating the need for the publisher to specify it explicitly.

See Also

- IceStorm Quality of Service
- Oneway Invocations
- Request Contexts
- How Glacier2 uses Request Contexts
- Batched Invocations
- Proxy Methods

Copyright 2016, ZeroC, Inc.

Implementing an IceStorm Subscriber

Our weather measurement subscriber implementation takes the following steps:

1. Obtain a proxy for the `TopicManager`. This is the primary IceStorm object, used by both publishers and subscribers.
2. Create an object adapter to host our `Monitor` servant.
3. Instantiate the `Monitor` servant and activate it with the object adapter.
4. Subscribe to the `Weather` topic.
5. Process `report` messages until shutdown.
6. Unsubscribe from the `Weather` topic.

We present monitor implementations in C++ and Java below.

On this page:

- Subscriber Example in C++
- Subscriber Example in Java

Subscriber Example in C++

Our C++ monitor implementation begins by including the necessary header files. The interesting ones are `IceStorm/IceStorm.h`, which is generated from the IceStorm Slice definitions, and `Monitor.h`, containing the generated code for our monitor definitions:

Copyright 2016, ZeroC, Inc.

C++

```
#include <Ice/Ice.h>
#include <IceStorm/IceStorm.h>
#include <Monitor.h>

using namespace std;

class MonitorI : virtual public Monitor {
public:
    virtual void report(const Measurement& m, const Ice::Current&) {
        cout << "Measurement report:" << endl
             << "  Tower: " << m.tower << endl
             << "  W Spd: " << m.windSpeed << endl
             << "  W Dir: " << m.windDirection << endl
             << "   Temp: " << m.temperature << endl
             << endl;
    }
};

int main(int argc, char* argv[])
{
    ...
    Ice::ObjectPrx obj = communicator->stringToProxy("IceStorm/TopicManager:tcp -p
9999");
    IceStorm::TopicManagerPrx topicManager =
IceStorm::TopicManagerPrx::checkedCast(obj);

    Ice::ObjectAdapterPtr adapter = communicator->createObjectAdapter("MonitorAdapt
er");

    MonitorPtr monitor = new MonitorI;
    Ice::ObjectPrx proxy = adapter->addWithUUID(monitor)->ice_oneway();
    adapter->activate();

    IceStorm::TopicPrx topic;
    try {
        topic = topicManager->retrieve("Weather");
        IceStorm::QoS qos;
        topic->subscribeAndGetPublisher(qos, proxy);
    }
    catch (const IceStorm::NoSuchTopic&) {
        // Error! No topic found!
        ...
    }

    communicator->waitForShutdown();

    topic->unsubscribe(proxy);
    ...
}
```


Copyright 2016, ZeroC, Inc.

Our implementation of the `Monitor` servant is currently quite simple. A real implementation might update a graphical display, or incorporate the measurements into an ongoing calculation.

C++

```
class MonitorI : virtual public Monitor {
public:
    virtual void report(const Measurement& m, const Ice::Current&) {
        cout << "Measurement report:" << endl
             << "  Tower: " << m.tower << endl
             << "  W Spd: " << m.windSpeed << endl
             << "  W Dir: " << m.windDirection << endl
             << "   Temp: " << m.temperature << endl
             << endl;
    }
};
```

After obtaining a proxy for the topic manager, the program creates an object adapter, instantiates the `Monitor` servant and activates it:

C++

```
    Ice::ObjectAdapterPtr adapter =
communicator->createObjectAdapter("MonitorAdapter");

    MonitorPtr monitor = new MonitorI;
    Ice::ObjectPrx proxy = adapter->addWithUUID(monitor)->ice_oneway();
    adapter->activate();
```

Note that the code creates a oneway proxy for the `Monitor` servant. This is for efficiency reasons: by subscribing with a oneway proxy, IceStorm will deliver events to the subscriber via oneway messages, instead of via twoway messages. We also activate the object adapter at this time, which means the servant can now begin receiving invocations.

Next, the monitor subscribes to the topic:

C++

```
    IceStorm::TopicPrx topic;
    try {
        topic = topicManager->retrieve("Weather");
        IceStorm::QoS qos;
        topic->subscribeAndGetPublisher(qos, proxy);
    }
    catch (const IceStorm::NoSuchTopic&) {
        // Error! No topic found!
        ...
    }
```

Finally, the monitor blocks until the communicator is shutdown. After `waitForShutdown` returns, the monitor cleans up by unsubscribing

Copyright 2016, ZeroC, Inc.

from the topic:

C++

```
adapter->activate();
communicator->waitForShutdown();

topic->unsubscribe(proxy);
```

Subscriber Example in Java

The Java implementation of the monitor is shown below:

Copyright 2016, ZeroC, Inc.

Java

```
class MonitorI extends _MonitorDisp {
    public void report(Measurement m, Ice.Current curr) {
        System.out.println(
            "Measurement report:\n" +
            "  Tower: " + m.tower + "\n" +
            "  W Spd: " + m.windSpeed + "\n" +
            "  W Dir: " + m.windDirection + "\n" +
            "   Temp: " + m.temperature + "\n");
    }
}

public static void main(String[] args)
{
    ...
    Ice.ObjectPrx obj = communicator.stringToProxy("IceStorm/TopicManager:tcp -p 99
99");
    IceStorm.TopicManagerPrx topicManager =
IceStorm.TopicManagerPrxHelper.checkedCast(obj);

    Ice.ObjectAdapterPtr adapter =
communicator.createObjectAdapter("MonitorAdapter");

    Monitor monitor = new MonitorI();
    Ice.ObjectPrx proxy = adapter.addWithUUID(monitor).ice_oneway();
    adapter.activate();

    IceStorm.TopicPrx topic = null;
    try {
        topic = topicManager.retrieve("Weather");
        java.util.Map qos = null;
        topic.subscribeAndGetPublisher(qos, proxy);
    }
    catch (IceStorm.NoSuchTopic ex) {
        // Error! No topic found!
        ...
    }

    communicator.waitForShutdown();

    topic.unsubscribe(proxy);
    ...
}
```

Our implementation of the `Monitor` servant is currently quite simple. A real implementation might update a graphical display, or incorporate the measurements into an ongoing calculation.

Copyright 2016, ZeroC, Inc.

Java

```
class MonitorI extends _MonitorDisp {
    public void report(Measurement m, Ice.Current curr) {
        System.out.println(
            "Measurement report:\n" +
            "  Tower: " + m.tower + "\n" +
            "  W Spd: " + m.windSpeed + "\n" +
            "  W Dir: " + m.windDirection + "\n" +
            "   Temp: " + m.temperature + "\n");
    }
}
```

After obtaining a proxy for the topic manager, the program creates an object adapter, instantiates the `Monitor` servant and activates it:

Java

```
    Monitor monitor = new MonitorI();
    Ice.ObjectPrx proxy = adapter.addWithUUID(monitor).ice_oneway();
    adapter.activate();
```

Note that the code creates a oneway proxy for the `Monitor` servant. This is for efficiency reasons: by subscribing with a oneway proxy, IceStorm will deliver events to the subscriber via oneway messages, instead of via twoway messages. We also activate the object adapter at this time, which means the servant can now begin receiving invocations.

Next, the monitor subscribes to the topic:

Java

```
    IceStorm.TopicPrx topic = null;
    try {
        topic = topicManager.retrieve("Weather");
        java.util.Map qos = null;
        topic.subscribeAndGetPublisher(qos, proxy);
    }
    catch (IceStorm.NoSuchTopic ex) {
        // Error! No topic found!
        ...
    }
```

Finally, the monitor blocks until the communicator is shutdown. After `waitForShutdown` returns, the monitor cleans up by unsubscribing from the topic:

Copyright 2016, ZeroC, Inc.

Java

```
communicator.waitForShutdown();

topic.unsubscribe(proxy);
```

See Also

- Using IceStorm
- Oneway Invocations

Copyright 2016, ZeroC, Inc.

Publishing to a Specific Subscriber

If you send events to the publisher object you obtain by calling `Topic::getPublisher`, the event is forwarded to all subscribers for that topic:

C++

```
IceStorm::TopicPrx topic = ...;
Ice::ObjectPrx pub = topic->getPublisher()->ice_oneway();

MonitorPrx monitor = MonitorPrx::uncheckedCast(pub);
Measurement m = ...;

monitor->report(m); // Sent to all subscribers
```

You can also publish an event to a single specific subscriber, by using the return value of `subscribeAndGetPublisher`. For example:

C++

```
MonitorPtr monitor = new MonitorI;
Ice::ObjectPrx proxy = adapter->addWithUUID(monitor)->ice_oneway();

IceStorm::topicPrx topic = ...;

Icestorm::QoS qos;
Ice::ObjectPrx pub = topic->subscribeAndGetPublisher(qos, proxy);
MonitorPrx monitor = MonitorPrx::uncheckedCast(pub);

Measurement m = ...;
monitor->report(m); // Sent only to this subscriber
```

Note that, here, we save the return value of `subscribeAndGetPublisher`. The return value is a proxy that connects specifically to the `MonitorI` instance denoted by `proxy`. However, when the code calls `report` on that proxy, instead of directly invoking on the `MonitorI` instance, the request is forwarded via IceStorm.

As it stands, this code is not very interesting. After all, the call to `monitor->report` is just a round-about way for the subscriber to publish a message to itself. However, the subscriber can pass this subscriber-specific publisher proxy to another process. When that process publishes an event via the proxy, the event is sent only to the specific subscriber, instead of to all subscribers for the topic. In turn, this is useful if you are using the observer pattern, with all observers attached to an IceStorm topic.

As an example, we might have a list whose state is to be monitored by a number of observers. Updates to the list are published to an IceStorm topic, say, `ListUpdates`. The observers of the list subscribe with an interface such as:

Slice

```
interface ListObserver {
    void init(/* The entire state of the list */);
    void itemChange(/* The added or deleted item */);
};
```


Copyright 2016, ZeroC, Inc.

The idea is that, when an observer first starts observing the list, the `init` operation is called on the observer and passed the entire list. This initializes the observer with the current state of the list. Thereafter, whenever the list changes, it calls `itemChange` on the observer to inform it of the addition or deletion of an item. (The details of how this happens are secondary; the important point is that the observer is informed of the current state of the list initially and, thereafter, receives incremental updates about modifications to the list, rather than the entire list whenever it changes.)

The list itself might look something like this:

Slice

```
interface List {
    void add(Item i);
    void remove(Item i);

    void addObserver(ListObserver* lo);
    void removeObserver(ListObserver* lo);
};
```

The list provides operations to add and remove an item, as well as operations to add and remove an observer. Every time `add` or `remove` ar e called on the list, the list publishes an `itemChange` event to the `ListUpdates` topic; this informs all the subscribed observers of the change to the list. However, when an observer is first added, the observer's `init` operation must be called. Moreover, we want to call that method only once for each observer, so we cannot just publish the initial state of the list on a topic that all observers subscribe to.

The subscriber-specific proxy that is returned by `subscribeAndGetPublisher` solves this nicely: the implementation of `addObserver` cal ls `subscribeAndGetPublisher`, and then invokes `init` on the observer. This both subscribes the observer to the topic, and IceStorm forwards the call to `init` to the observer. This is preferable to the list invoking `init` on the observer directly: if the observer is misbehaved (for example, if its `init` implementation blocks for some time), the list is unaffected because IceStorm shields the list from such behavior.
See Also

- Using IceStorm

Copyright 2016, ZeroC, Inc.

Highly Available IceStorm

IceStorm offers a highly available (HA) mode that employs master-slave replication with automatic failover in case the master fails.

On this page:

- IceStorm Replication Algorithm
- IceStorm Replica States
- Client Considerations for IceStorm Replication
- Subscriber Considerations for IceStorm Replication
- Publisher Considerations for IceStorm Replication

IceStorm Replication Algorithm

HA IceStorm uses the Garcia-Molina "Invitation Election Algorithm" [1] in which each replica has a priority and belongs to a replica group. The replica with the highest priority in the group becomes the coordinator, and the remaining replicas are slaves of the coordinator.

All replicas are statically configured with information about all other replicas, including their priority. The group combining works as follows:

- When recovering from an error, or during startup, replicas form a single self-coordinated group.
- Coordinators periodically attempt to combine their groups with other groups in order to form larger groups.

At regular intervals, slave replicas contact their coordinator to ensure that the coordinator is still the master of the slave's group. If a failure occurs, the replica considers itself in error and performs error recovery as described above.

Replication commences once a group contains a majority of replicas. A majority is necessary to avoid the possibility of network partitioning, in which two groups of replicas form that cannot communicate and whose database contents diverge. With respect to IceStorm, a consequence of requiring a majority is that a minimum of three replicas are necessary.

An exception to the majority rule is made during full system startup (i.e., when no replica is currently running). In this situation, replication can only commence with the participation of every replica in the group. This requirement guarantees that the databases of all replicas are synchronized, and avoids the risk that the database of an offline replica might contain more recent information.

Once a majority group has been formed, all database states are compared. The most recent database state (as determined by comparing a time stamp recorded upon each database change) is transferred to all replicas and replication commences. IceStorm is now available for use.

IceStorm Replica States

IceStorm replicas can have one of four states:

- Inactive:
  The node is inactive and awaiting an election.
- Election:
  The node is electing a coordinator.
- Reorganization:
  The replica group is reorganizing.
- Normal:
  The replica group is active and replicating.

For debugging purposes, you can obtain the state of the replicas using the `replica` command of the `icestormadmin` utility, as shown below:

Copyright 2016, ZeroC, Inc.

```
$ icestormadmin --Ice.Config=config
>>> replica
replica count: 3
1: id:         1
1: coord:      3
1: group name: 3:191131CC-703A-41D6-8B80-D19F0D5F0410
1: state:      normal
1: group:
1: max:        3
2: id:         2
2: coord:      3
2: group name: 3:191131CC-703A-41D6-8B80-D19F0D5F0410
2: state:      normal
2: group:
2: max:        3
3: id:         3
3: coord:      3
3: group name: 3:191131CC-703A-41D6-8B80-D19F0D5F0410
3: state:      normal
3: group:      1,2
3: max:        3
```

Each line begins with the identifier of the replica. The command displays the following information:

- `id`
  The identifier of the replica.
- `coord`
  The identifier of the group's coordinator.
- `group name`
  The name of the group to which this replica belongs.
- `state`
  The replica's current state.
- `group`
  The identifiers of the other replicas in the group. Note that only the coordinator knows, or cares about, this information.
- `max`
  The maximum number of replicas seen by this replica. This value is used during startup to determine whether full participation is necessary. If the value is less than the total number of replicas, full participation is required.

Client Considerations for IceStorm Replication

As previously noted, an individual IceStorm replica can be in one of several states. However, IceStorm clients have a different perspective in which the replication group as a whole is in one of the states shown below:

- Down
  All requests to IceStorm fail.
- Inactive
  All requests to IceStorm block until the replica is either down (in which case the request fails), or becomes Active.
- Active
  Requests are processed.

It is also possible, but highly unlikely, for a request to result in an `Ice::UnknownException`. This can happen, for example, if a replica loses the majority and thus progresses to the inactive state during request processing. In this case, the result of the request is indeterminate (the request may or may not have succeeded) and therefore the IceStorm client can draw no conclusion. The client should retry the request

Copyright 2016, ZeroC, Inc.

and be prepared for the request to fail. Consider this example:

C++

```
TopicPrx topic = ...;
Ice::ObjectPrx sub = ...;
IceStorm::QoS qos;
topic->subscribeAndGetPublisher(qos, sub);
```

The call to `subscribeAndGetPublisher` may fail in very rare cases with an `UnknownException`, indicating that the subscription may or may not have succeeded. Here is the proper way to deal with the possibility of an `UnknownException`:

C++

```
TopicPrx topic = ...;
Ice::ObjectPrx sub = ...;
IceStorm::QoS qos;
while(true) {
    try {
        topic->subscriberAndGetPublisher(qos, sub);
    } catch(const Ice::UnknownException&) {
        continue;
    } catch(const IceStorm::AlreadySubscribed&) {
        // Expected.
    }
    break;
}
```

Subscriber Considerations for IceStorm Replication

Subscribers can receive events from any replica. The subscriber will stop receiving events under two circumstances:

- The subscriber is unsubscribed by calling `Topic::unsubscribe`.
- The subscriber is removed as a result of a failure to deliver events. Subscribers can optionally configure a quality of service paramet er that causes IceStorm to make additional delivery attempts.

Publisher Considerations for IceStorm Replication

A publisher for HA IceStorm typically receives a proxy containing multiple endpoints. With this proxy, the publisher normally binds to a single replica and continues using that replica unless there is a failure, or until active connection management (ACM) closes the connection.

As with non-HA IceStorm, event delivery ordering can be guaranteed if the subscriber and publisher are suitably configured and the publisher continues to use the same replica when publishing events.

Ordering guarantees are lost as soon as a publisher changes to a different replica. Furthermore, a publisher may receive no notification that a change has occurred, which is possible under two circumstances:

- ACM has closed the connection.
- Publishing to a replica fails and the Ice invocation can be retried, in which case the Ice run time in the publisher automatically and transparently attempts to send the request to another replica. The publisher receives an exception if the invocation cannot be retried.

A publisher has two ways of ensuring that it is notified about a change in replicas:

- The simplest method is to use the `Topic::getNonReplicatedPublisher` operation. The proxy returned by this operation points

Copyright 2016, ZeroC, Inc.

directly at the current replica and no transparent failover to a different can occur.
- If you never want transparent failover to occur during publishing, you can configure your publisher proxy so that it contains only one endpoint. In this configuration, the `Topic::getPublisher` operation behaves exactly like `getNonReplicatedPublisher`.

Of the two strategies, using `getNonReplicatedPublisher` is preferable for two reasons:

- It does not involve changes to IceStorm's configuration.
- It is still possible to obtain a replicated publisher proxy by calling `getPublisher`, whereas if you had used the second strategy you would have eliminated that possibility.

The second strategy may be necessary in certain circumstances, such as when an existing IceStorm application is deployed and cannot be changed.

Regardless of the strategy you choose, a publisher can recover from the failure of a replica by requesting another proxy from the replicated topic using `getPublisher` or `getNonReplicatedPublisher`.

See Also

- IceStorm Administration
- IceStorm Quality of Service
- Active Connection Management
- IceStorm Delivery Modes
- Automatic Retries
- Configuring IceStorm

References

1. Garcia-Molina, H. 1982. Elections in a Distributed Computing System. *IEEE Transactions on Computers* 31 (1): 48-59.

Copyright 2016, ZeroC, Inc.

IceStorm Administration

The IceStorm administration tool, `icestormadmin`, is a command-line program that provides administrative control of an IceStorm server.

Configuration properties inform `icestormadmin` about the topic manager(s) that you wish to administer. You have several configuration options:

- Define `IceStormAdmin.Host` and `IceStormAdmin.Port` - `icestormadmin` constructs its own proxy for the topic manager at the specified host and port
- Define `IceStormAdmin.TopicManager.Default` - specifies a proxy for the default topic manager on which administrative commands operate
- Define `IceStormAdmin.TopicManager.`*`name`* - specifies the proxies for any number of named topic managers

The tool supports the following command-line options:

```
$ icestormadmin -h
Usage: icestormadmin [options]
Options:
-h, --help           Show this message.
-v, --version        Display the Ice version.
-e COMMANDS          Execute COMMANDS.
-d, --debug          Print debug messages.
```

If you specify one or more `-e` options, the tool executes the given commands and exits, otherwise the tool enters an interactive session. The `help` command displays the following usage information:

- `help`
  Print this message.
- `exit, quit`
  Exit this program.
- `create` *`TOPICS`*
  Add *`TOPICS`*.
- `destroy` *`TOPICS`*
  Remove *`TOPICS`*.
- `link` *`FROM TO`* [*`COST`*]
  Link *`FROM`* to *`TO`* with the optional *`COST`*.
- `unlink` *`FROM TO`*
  Unlink *`TO`* from *`FROM`*.
- `links` [*`INSTANCE-NAME`*]
  Without an argument, `links` displays the links of all topics in the current topic manager. You can specify a different topic manager by providing its instance name.
- `topics` [*`INSTANCE-NAME`*]
  Without an argument, *`topics`* displays the names of all topics in the current topic manager. You can specify a different topic manager by providing its instance name.
- `current` [*`INSTANCE-NAME`*]
  Set the current topic manager to the topic manager with instance name *`INSTANCE-NAME`*. The proxy of the corresponding topic manager must be specified by setting an `IceStormAdmin.TopicManager.`*`name`* property. Without an argument, the command shows the current topic manager.
- `replica` [*`INSTANCE-NAME`*]
  Display replication information for the given *`INSTANCE-NAME`*.
- `subscribers` *`TOPICS`*
  Displays the identities of the subscribers for each of the *`TOPICS`*.

Some of the commands accept one or more topic names (*`TOPICS`*) as arguments. Topic names containing white space or matching a

Copyright 2016, ZeroC, Inc.

command keyword must be enclosed in single or double quotes.

By default, `icestormadmin` uses the topic manager specified by your setting for `IceStormAdmin.TopicManager.Default`. For example, without additional arguments, the `create` command operates on that topic manager.

If you are using multiple topic managers, you can specify their proxies by setting `IceStormAdmin.TopicManager.`*`name`* for each topic manager. For example:

```
IceStormAdmin.TopicManager.A=A/TopicManager:tcp -h x -p 9995
IceStormAdmin.TopicManager.B=Foo/TopicManager:tcp -h x -p 9996
IceStormAdmin.TopicManager.C=Bar/TopicManager:tcp -h z -p 9995
```

This sets the proxies for three topic managers. Note that *name* need not match the instance name of the corresponding topic manager — *name* simply serves as a tag. With these property settings, the `icestormadmin` commands that accept a topic can now specify a topic manager other than the default topic manager that is configured with `IceStormAdmin.TopicManager.Default`. For example:

```
current Foo
create myTopic
create Bar/myOtherTopic
```

This sets the current topic manager to the one with instance name `Foo`; the first `create` command then creates the topic within that topic manager, whereas the second `create` command uses the topic manager with instance name `Bar`.

See Also

- Configuring IceStorm
- Highly Available IceStorm
- IceStorm Properties
- IceStormAdmin.*

Copyright 2016, ZeroC, Inc.

Topic Federation

The ability to link topics together into a federation provides IceStorm applications with a lot of flexibility, while the notion of a "cost" associated with links allows applications to restrict the flow of messages in creative ways. IceStorm applications have complete control of topic federation using the `TopicManager` interface described in the online XREF Slice API Reference, allowing links to be created and removed dynamically as necessary. For many applications, however, the topic graph is static and therefore can be configured using the administrative tool .

On this page:

- IceStorm Message Propagation
- Using Cost to Limit Message Propagation
  - Request Context for Cost
  - Publishing a Message with a Cost
  - Receiving a Message with a Cost
- Automating IceStorm Federation
  - Administration Tool Script
- Proxy Considerations for IceStorm Federation

IceStorm Message Propagation

IceStorm messages are never propagated over more than one link. For example, consider the topic graph shown below:

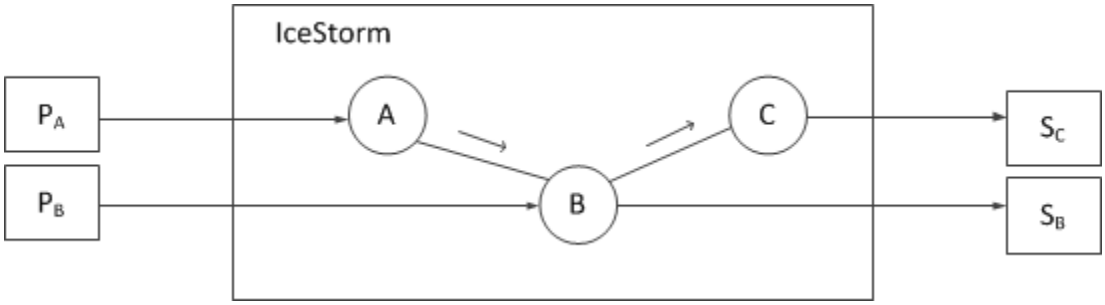

Message propagation.

In this case, messages published on `A` are propagated to `B`, but `B` does not propagate `A`'s messages to `C`. Therefore, subscriber S_B receives messages published on topics `A` and `B`, but subscriber S_C only receives messages published on topics `B` and `C`. If the application needs messages to propagate from `A` to `C`, then a link must be established directly between `A` and `C`.

Using Cost to Limit Message Propagation

As described above, IceStorm messages are only propagated on the originating topic's immediate links. In addition, applications can use the notion of cost to further restrict message propagation.

A cost is associated with messages and links. When a message is published on a topic, the topic compares the cost associated with each of its links against the message cost, and only propagates the message on those links whose cost equals or exceeds the message cost. A cost value of zero (`0`) has the following implications:

- messages with a cost value of zero (`0`) are published on all of the topic's links regardless of the link cost;
- links with a cost value of zero (`0`) accept all messages regardless of the message cost.

For example, consider the following topic graph:

Copyright 2016, ZeroC, Inc.

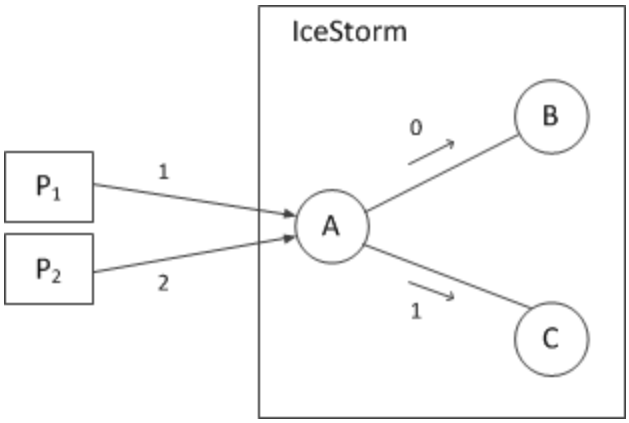

Cost semantics.

Publisher P_1 publishes a message on topic `A` with a cost of `1`. This message is propagated on the link to topic `B` because the link has a cost of `0` and therefore accepts all messages. The message is also propagated on the link to topic `C`, because the message cost does not exceed the link cost (`1`). On the other hand, the message published by P_2 with a cost of `2` is only propagated on the link to `B`.

Request Context for Cost

The cost of a message is specified in an Ice request context. Each Ice proxy operation has an implicit argument of type `Ice::Context` representing the request context. This argument is rarely used, but it is the ideal location for specifying the cost of an IceStorm message because an application only needs to supply a request context if it actually uses IceStorm's cost feature. If the request context does not contain a cost value, the message is assigned the default cost value of zero (0).

Publishing a Message with a Cost

The code examples below demonstrate how a collector can publish a measurement with a cost value of `5`. First, the C++ version:

C++

```
Measurement m = getMeasurement();
Ice::Context ctx;
ctx["cost"] = "5";
monitor->report(m, ctx);
```

And here is the equivalent version in Java:

Java

```
Measurement m = getMeasurement();
java.util.HashMap ctx = new java.util.HashMap();
ctx.put("cost", "5");
monitor.report(m, ctx);
```

Receiving a Message with a Cost

A subscriber can discover the cost of a message by examining the request context supplied in the `Ice::Current` argument. For example, here is a C++ implementation of `Monitor::report` that displays the cost value if it is present:

Copyright 2016, ZeroC, Inc.

C++

```
virtual void report(const Measurement& m, const Ice::Current& curr) {
    Ice::Context::const_iterator p = curr.ctx.find("cost");
    cout << "Measurement report:" << endl
         << "  Tower: " << m.tower << endl
         << "  W Spd: " << m.windSpeed << endl
         << "  W Dir: " << m.windDirection << endl
         << "   Temp: " << m.temperature << endl
         << "   Temp: " << m.temperature << endl;
    if (p != curr.ctx.end())
        cout << "   Cost: " << p->second << endl;
    cout << endl;
}
```

And here is the equivalent Java implementation:

Java

```
public void report(Measurement m, Ice.Current curr) {
    String cost = null;
    if (curr.ctx != null)
        cost = curr.ctx.get("cost");
    System.out.println(
        "Measurement report:\n" +
        "  Tower: " + m.tower + "\n" +
        "  W Spd: " + m.windSpeed + "\n" +
        "  W Dir: " + m.windDirection + "\n" +
        "   Temp: " + m.temperature);
    if (cost != null)
        System.out.println("   Cost: " + cost);
    System.out.println();
}
```

For the sake of efficiency, the Ice for Java run time may supply a null value for the request context in `Ice.Current`, therefore an application is required to check for null before using the request context.

Automating IceStorm Federation

Given the restrictions on message propagation described in the previous sections, creating a complex topic graph can be a tedious endeavor. Of course, creating a topic graph is not typically a common occurrence, since IceStorm keeps a persistent record of the graph. However, there are situations where an automated procedure for creating a topic graph can be valuable, such as during development when the graph might change significantly and often, or when graphs need to be recomputed based on changing costs.

Administration Tool Script

A simple way to automate the creation of a topic graph is to create a text file containing commands to be executed by the IceStorm administration tool. For example, the commands to create the topic graph shown earlier are shown below:

Copyright 2016, ZeroC, Inc.

```
create A B C
link A B 0
link A C 1
```

If we store these commands in the file `graph.txt`, we can execute them using the following command:

```
$ icestormadmin --Ice.Config=config < graph.txt
```

We assume that the configuration file `config` contains the definition for the property `IceStormAdmin.TopicManager.Default`.

Proxy Considerations for IceStorm Federation

Note that, if you federate IceStorm servers, you must ensure that the proxies for the linked topics always use the same host and port (or, alternatively, can be indirectly bound via IceGrid), otherwise the federation cannot be re-established if one of the servers in the federation shuts down and is restarted later.

See Also

- IceStorm Administration
- Request Contexts
- The Current Object
- IceStorm Properties
- IceGrid

Copyright 2016, ZeroC, Inc.

IceStorm Quality of Service

An IceStorm subscriber specifies Quality of Service (QoS) parameters at the time of subscription. The supported QoS parameters are described in the sections below

On this page:

- Reliability QoS for IceStorm
- Retry Count QoS for IceStorm
- IceStorm QoS Example

Reliability QoS for IceStorm

The QoS parameter `reliability` affects message delivery. The only legal values at this point are `ordered` and the empty string. If not specified, the default value is the empty string (meaning not ordered).

The `ordered` reliability QoS requires a twoway subscriber proxy. If you specify this reliability QoS, IceStorm will forward events in the order they are received but doesn't forward them immediately. Instead, IceStorm waits for the reply from the forwarding of an event before forwarding the next event. This guarantees that the subscriber will process the events in the same order as they were received even if its thr eading model doesn't serialize incoming requests.

Retry Count QoS for IceStorm

IceStorm automatically removes a subscriber if `ObjectNotExistException` or `NotRegisteredException` is raised while attempting to deliver an event. IceStorm considers these exceptions as indicators of a hard failure, after which it is unnecessary to continue event delivery.

For other kinds of failures, IceStorm uses the QoS parameter `retryCount` to determine when to remove a subscriber. A value of `-1` means IceStorm retries forever and never automatically removes a subscriber unless a hard failure occurs. A value of zero means IceStorm never retries and immediately removes the subscriber. For positive values, IceStorm decrements the subscriber's retry count for each failure and removes the subscriber once it reaches zero. Linked topics always have a configured retry count of `-1`. The default value of the `retryCoun t` parameter is zero.

A retry count of `-1` adds some resiliency to your IceStorm application by ignoring intermittent network failures such as `ConnectionRefuse dException`. However, there is also some risk inherent in using a retry count of `-1` because an improperly configured subscriber may never be removed. For example, consider what happens when a subscriber registers using a transient endpoint: if that subscriber happens to terminate and resubscribe with a different endpoint, IceStorm will continue trying to deliver events to the subscriber at its old endpoint. IceStorm can only remove the subscriber if it receives a hard error, and that is only possible when the subscriber is reachable.

To use a retry count of `-1` successfully, the subscriber should either register with a fixed endpoint, or use IceGrid to take advantage of indirect proxies and automatic activation. Furthermore, if the subscriber is expected to function correctly after a restart of its process, the subscriber must use the same identity. The application can rely on the `subscribeAndGetPublisher` operation to raise `AlreadySubscr ibed` when the subscriber is already subscribed.

IceStorm QoS Example

The Slice type `IceStorm::QoS` is defined as a `dictionary` whose key and value types are both `string`, therefore the QoS parameter name and value are both represented as strings. The code we presented in our earlier subscriber example used an empty dictionary for the QoS argument, meaning default values are used. The C++ and Java examples shown below illustrate how to set the `reliability` paramet er to `ordered`.

Here is the C++ example:

C++

```
IceStorm::QoS qos;
qos["reliability"] = "ordered";
topic->subscribeAndGetPublisher(qos, proxy->ice_twoway());
```

Here is the Java example:

Copyright 2016, ZeroC, Inc.

Java

```
java.util.Map qos = new java.util.HashMap();
qos.put("reliability", "ordered");
topic.subscribeAndGetPublisher(qos, proxy.ice_twoway());
```

See Also

- IceGrid
- Implementing an IceStorm Subscriber
- IceStorm Delivery Modes
- Object Identity

Copyright 2016, ZeroC, Inc.

IceStorm Delivery Modes

The delivery mode for events sent to subscribers is controlled by the proxy that the subscriber passes to IceStorm. For example, if the subscriber subscribes with a oneway proxy, events will be forwarded by IceStorm as oneway messages.

On this page:

- Subscribing with a Twoway Proxy
- Subscribing with a Oneway Proxy
- Subscribing with a Batch Oneway Proxy
- Subscribing with a Datagram Proxy
- Subscribing with a Batch Datagram Proxy

Subscribing with a Twoway Proxy

In this mode each event is sent to the subscriber as a separate twoway message. This allows the subscriber to enable server-side active connection management (ACM) without risking lost messages, because IceStorm will re-send an event if the subscriber happens to close its connection at the wrong moment.

If you combine a twoway proxy with a `reliability` QoS of `ordered`, messages will be forwarded to the subscriber in the order in which they are received. This is guaranteed because IceStorm will wait for a reply from the subscriber for each event before sending the next event.

Without ordered delivery, events may be delivered out-of-order to the subscriber because IceStorm will send an event as soon as possible (without waiting for a reply for the preceding event). If the subscriber uses a thread pool with more than one thread, this can result in out-of-order dispatch of messages in the subscriber.

For single-threaded subscribers and subscribers using a serialized thread pool, twoway delivery always produces in-order dispatch of events in the subscriber.

With twoway delivery, IceStorm is informed of any failure to deliver an event by the Ice run time. For example, IceStorm may not be able to establish a connection to a subscriber, or may receive an `ObjectNotExistException` when it forwards an event. Any failure to deliver an event to a subscriber (possibly after a transparent retry by the Ice run time) results in the cancellation of the corresponding subscription.

Subscribing with a Oneway Proxy

In this mode each event is sent to the subscriber as a oneway message. If more than one event is ready to be delivered, the events are sent in a single batch. This delivery mode is more efficient than using twoway delivery. However, the subscriber cannot use active connection management without the risk of events being lost. In addition, if something goes wrong with the subscriber, such as the subscriber having destroyed its callback object without unsubscribing, or having subscribed an object with the wrong interface, IceStorm does not notice the failure and will continue to send events to the non-existent subscriber object for as long as it can maintain a connection to the subscriber's endpoint.

For multi-threaded subscribers, oneway delivery can result in out-of-order delivery of events. For single-threaded subscribers and subscribers using a serialized thread pool, events are delivered in order.

Subscribing with a Batch Oneway Proxy

With this delivery mode, IceStorm buffers events from publishers and sends them in batches to the subscriber. This reduces network overhead and is more efficient than oneway delivery. However, as for oneway delivery, the subscriber cannot use active connection management without the risk of losing events. In addition, events can be delivered out of order if the subscriber is multi-threaded. Batch oneway delivery, while providing better throughput, increases latency because events arrive in "bursts". You can control the interval at which batched events are flushed by setting the `IceStorm.Flush.Timeout` property.

Subscribing with a Datagram Proxy

With this delivery mode, events are forwarded as UDP messages, optionally with multicast semantics. This means that events can be delivered out of order, can be lost, and can even be duplicated. In addition, IceStorm cannot detect anything about the delivery status of events. This means that if a subscriber disappears without unsubscribing, IceStorm will attempt to forward events to the subscriber indefinitely. If you use datagram delivery, you need to be careful that subscribers unsubscribe before they disappear; otherwise, stale subscriptions can accumulate in IceStorm over time, bogging down the service as it delivers more and more events to subscribers that no longer exist.

Copyright 2016, ZeroC, Inc.

Subscribing with a Batch Datagram Proxy

With this delivery mode, events are forwarded as batches within a datagram. The same considerations as for datagram delivery and oneway batched delivery apply here. In addition, keep in mind that, due to the size limit for datagrams, batched datagram delivery makes sense only if events are small. (You should also consider enabling compression with this delivery mode.)

See Also

- Active Connection Management
- Oneway Invocations
- The Ice Threading Model
- Batched Invocations
- IceStorm Properties

Copyright 2016, ZeroC, Inc.

Configuring IceStorm

IceStorm is a relatively lightweight service in that it requires very little configuration and is implemented as an IceBox service. The configuration properties supported by IceStorm are described in IceStorm Properties; some of them control diagnostic output and are not discussed here.

On this page:

- IceStorm Property Prefix
- IceStorm Server Configuration
- Deploying IceStorm Replicas
  - IceGrid Deployment
  - Manual Deployment
- IceStorm Client Configuration
- IceStorm Object Identities
- Using the IceStorm Finder Interface

IceStorm Property Prefix

As you will see in IceStorm Properties, IceStorm uses its IceBox service name as the prefix for all of its properties. For example, the property `service.TopicManager.Endpoints` becomes `DemoIceStorm.TopicManager.Endpoints` when IceStorm is configured as the IceBox service `DemoIceStorm`.

IceStorm Server Configuration

The first step is configuring IceBox to run the IceStorm service:

```
IceBox.Service.DemoIceStorm=IceStormService,34:createIceStorm --Ice.Config=config.s
ervice
```

In this example, the IceStorm service itself is configured by the properties in the `config.service` file, which might look as follows for a non-replicated service:

```
Freeze.DbEnv.DemoIceStorm.DbHome=db
DemoIceStorm.TopicManager.Endpoints=tcp -p 9999
DemoIceStorm.Publish.Endpoints=tcp -p 10000
```

IceStorm uses Freeze to manage the service's persistent state, therefore the first property specifies the path name of the Freeze database environment directory for the service. Here the directory `db` is used, which must already exist in the current working directory. This property can be omitted when the service is running in transient mode.

The final two properties specify the endpoints used by the IceStorm object adapters; notice that their property names begin with `DemoIceStorm`, matching the service name. The `TopicManager` property specifies the endpoints on which the `TopicManager` and `Topic` objects reside; these endpoints must use a connection-oriented protocol such as TCP or SSL. The `Publish` property specifies the endpoints used by topic publisher objects; using datagram endpoints in this property is possible but carries additional risk.

IceStorm's default thread pool configuration is sufficient when the service is running on a single CPU machine. On a host with multiple CPUs, you may be able to improve IceStorm's performance by increasing the size of its client-side thread pool using the `Ice.ThreadPool.Client.*` properties, but the optimal number of threads can only be determined with careful benchmarking.

Deploying IceStorm Replicas

There are two ways of deploying IceStorm in its highly available (replicated) mode. In both cases, adding another replica requires that all active replicas be stopped while their configurations are updated; it is not possible to add a replica while replication is running.

To remove a replica, stop all replicas and alter the configuration as necessary. You must be careful not to remove a replica if it has the latest database state. This situation will never occur during normal operation since the database state of all replicas is identical. However, in the event of a crash it is possible for a coordinator to have later database state than all replicas. The safest approach is to verify that all replicas are active prior to stopping them. You can do this using the `icestormadmin` utility by checking that all replicas are in the `Normal` state.

Copyright 2016, ZeroC, Inc.

IceGrid Deployment

IceGrid is a convenient way of deploying IceStorm replicas. The term *replica* is also used in the context of IceGrid, specifically when referring to groups of object adapters that participate in replication. It is important to be aware of the distinction between IceStorm replication and object adapter replication; IceStorm replication *uses* object adapter replication when deployed with IceGrid, but IceStorm does not *require* ob ject adapter replication as you will see below.

An IceGrid deployment typically uses two adapter replica groups: one for the publisher proxies, and another for the topics, as shown below:

XML

```
<replica-group id="DemoIceStorm-PublishReplicaGroup">
</replica-group>

<replica-group id="DemoIceStorm-TopicManagerReplicaGroup">
    <object identity="DemoIceStorm/TopicManager" type="::IceStorm::TopicManager"/>
</replica-group>
```

The object adapters are then configured to use these replica groups:

XML

```
<adapter name="${service}.Publish"
    endpoints="tcp"
    replica-group="${instance-name}-PublishReplicaGroup"/>

<adapter name="${service}.TopicManager"
    endpoints="tcp"
    replica-group="${instance-name}-TopicManagerReplicaGroup"/>
```

An application may not want publisher proxies to contain multiple endpoints. In this case you should remove `PublishReplicaGroup` from the above deployment.

The next step is defining the endpoints for the adapter `Node`, which is used internally for communication with other IceStorm replicas and is not part of an adapter replica group:

XML

```
<adapter name="${service}.Node" endpoints="tcp"/>
```

Finally, you must define the node ID for each IceStorm replica using the `NodeId` property. The node ID must be a non-negative integer:

XML

```
<property name="${service}.NodeId" value="${index}"/>
```

Example
You can find a complete C++ example of an IceGrid deployment in the directory `demo/IceStorm/replicated`.

Manual Deployment

Copyright 2016, ZeroC, Inc.

You can also deploy IceStorm replicas without IceGrid, although it requires more manual configuration; an IceGrid deployment is simpler to maintain.

The first step is defining the set of node proxies using properties of the form `Nodes.id`. These proxies allow replicas to contact each other; their object identities are composed using *`instance-name`*`/node`*`id`*.

For example, assuming we are using the IceBox service name `IceStorm` and have three replicas with the identifiers 0, 1, 2 and an instance name of `DemoIceStorm`, we can configure the proxies as shown below:

```
IceStorm.InstanceName=DemoIceStorm
IceStorm.Nodes.0=DemoIceStorm/node0:tcp -p 13000
IceStorm.Nodes.1=DemoIceStorm/node1:tcp -p 13010
IceStorm.Nodes.2=DemoIceStorm/node2:tcp -p 13020
```

These properties must be defined in each replica. Additionally, each replica must define its node ID, as well as the node's endpoints. For example, we can configure node 0 as follows:

```
IceStorm.NodeId=0
IceStorm.Node.Endpoints=tcp -p 13000
```

The endpoints for each replica and ID must match the proxies configured in the `Nodes.`*`id`* properties.

Two additional properties allow you to configure replicated endpoints:

- *`service-name`*`.ReplicatedTopicManagerEndpoints`
  Defines the endpoints contained in proxies returned by the topic manager.
- *`service-name`*`.ReplicatedPublishEndpoints`
  Defines the endpoints contained in the publisher proxy returned by the topic.

For example, suppose we configure three replicas:

```
IceStorm.NodeId=0
IceStorm.TopicManager.Endpoints=tcp -p 10000
IceStorm.Publish.Endpoints=tcp -p 10001:udp -p 10001

IceStorm.NodeId=1
IceStorm.TopicManager.Endpoints=tcp -p 10010
IceStorm.Publish.Endpoints=tcp -p 10011:udp -p 10011

IceStorm.NodeId=2
IceStorm.TopicManager.Endpoints=tcp -p 10020
IceStorm.Publish.Endpoints=tcp -p 10021:udp -p 10021
```

Each replica should also define these properties:

```
IceStorm.ReplicatedPublishEndpoints=\
    tcp -p 10001:tcp -p 10011:tcp -p 10021:udp -p 10001:udp -p 10011:udp -p 10021
IceStorm.ReplicatedTopicManagerEndpoints=tcp -p 10000:tcp -p 10010:tcp -p 10020
```

An application may not want publisher proxies to contain multiple endpoints. In this case you should remove the definition of the `ReplicatedPublishEndpoints` property from the above deployment.

Example

Copyright 2016, ZeroC, Inc.

You can find a complete C++ example of a manual deployment in the directory `demo/IceStorm/replicated2`.

IceStorm Client Configuration

Clients of the service can define a proxy for the `TopicManager` object as follows:

```
TopicManager.Proxy=IceStorm/TopicManager:tcp -p 9999
```

The name of the property is not relevant, but the endpoint must match that of the *`service`*`.TopicManager.Endpoints` property, and the object identity must use the IceStorm instance name as the category and `TopicManager` as the name.

IceStorm Object Identities

IceStorm hosts a well-known object that implements the `IceStorm::TopicManager` interface. The default identity of this object is `IceStorm/TopicManager`, as seen in the stringified proxy example above. If an application requires the use of multiple IceStorm services, it's a good idea to assign unique identities to their well-known objects by configuring the services with different values for the *`service`*`.InstanceName` property, as shown in the following example:

```
DemoIceStorm.InstanceName=Measurement
```

This property changes the category of the object's identity, which becomes `Measurement/TopicManager`. The client's configuration must also be changed to reflect the new identity:

```
TopicManager.Proxy=Measurement/TopicManager:tcp -p 9999
```

As of Ice 3.6, IceStorm also hosts an object with the identity `IceStorm/Finder`, as described in the next section. This identity is not affected by the service name or by changes to *`service`*`.InstanceName`.

Using the IceStorm `Finder` Interface

As of Ice 3.6, IceStorm supports the `IceStorm::Finder` interface:

Slice

```
module IceStorm {
interface Finder {
    TopicManager* getTopicManager();
};
};
```

An object supporting this interface is available with the identity `IceStorm/Finder` on the service's topic manager endpoint. By knowing the host and port of this endpoint, a client can discover the topic manager's proxy at run time with a call to `getTopicManager`:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPrx prx = communicator->stringToProxy("IceStorm/Finder:tcp -p 8888 -h
icestormhost");
IceStorm::FinderPrx finder = IceStorm::FinderPrx::checkedCast(prx);
IceStorm::TopicManagerPrx topicManager = finder->getTopicManager();
```

See Also

- IceStorm Properties
- IceBox
- Freeze
- IceGrid
- The Ice Threading Model
- Object Adapter Replication
- IceStorm Administration
- Using an IceStorm Publisher Object
- Highly Available IceStorm

Copyright 2016, ZeroC, Inc.

IceStorm Metrics

You can monitor IceStorm using the Administrative Facility and the Metrics Facet. IceStorm provides two metrics class to monitor topic and subscriber related metrics. These classes are defined in `IceStorm/Metrics.ice` and are shown below.

Slice

```
namespace IceMX {
class TopicMetrics extends Metrics {
    long published = 0;
    long forwarded = 0;
};
class SubscriberMetrics extends Metrics {
    int queued = 0;
    int outstanding = 0;
    long delivered = 0;
};
};
```

IceStorm records metrics in the metrics map described below.

Metrics map name	Slice class	Description	Property prefix
Topic	`IceMX::TopicMetrics`	Topic metrics	`IceMX.Metrics.`*`view-name`*`.Map.Topic`
Subscriber	`IceMX::SubscriberMetrics`	Subscriber metrics	`IceMX.Metrics.`*`view-name`*`.Map.Subscriber`

To configure a metrics view to record IceStorm topic and subscriber you can use for example:

- `IceMX.Metrics.IceStormView.Map.Topic.GroupBy=id`
- `IceMX.Metrics.IceStormView.Map.Subscriber.GroupBy=id`

This will configure a view containing only the `Topic` and `Subscriber` maps. All the topics and subscribers from the IceStorm service will be monitored individually with separate metrics object.

You can use the following attributes when configuring the IceStorm `Topic` map:

Name	Description
id	The id is the topic name.
parent	The IceStorm service name.
none	The empty string.
topic	The topic name.
service	The IceStorm service name.

The `Subscriber` map can be configured with the following attributes:

Name	Description
id	The id of the subscriber metrics is the stringified proxy of the subscriber.
parent	The name of the topic name to which this subscriber belongs.
none	The empty string.
topic	The name of the topic name to which this subscriber belongs.
service	The IceStorm service name.

Copyright 2016, ZeroC, Inc.

identity	The identity of the subscriber proxy.
facet	The facet of the subscriber proxy.
encoding	The encoding of the subscriber proxy.
mode	The mode of the subscriber proxy.
proxy	The subscriber proxy.
link	The proxy of the topic linked to the the topic which owns this subscriber.
state	The state of the subscriber. It can either be "online", "offline" or "error".

See Also

- Administrative Facility
- The Metrics Facet
- IceMX.Metrics.*

Copyright 2016, ZeroC, Inc.

IceStorm Database Utility

The `icestormdb` utility is a command-line tool for importing and exporting an IceStorm database. This tool is especially useful when migrating to Ice 3.7, where IceStorm uses a different database implementation and storage format. Using `icestormdb`, you can export your existing IceStorm 3.6 database and then import the data using the Ice 3.7 version of `icestormdb`.

On this page:

- Usage
- Exporting an IceStorm Database
- Importing an IceStorm Database

Usage

The IceStorm Database utility supports the following command-line options:

```
Usage: icestormdb <options>
Options:
-h, --help            Show this message.
-v, --version         Display version.
--import FILE         Import database from FILE.
--export FILE         Export database to FILE.
--dbhome DIR          The database directory.
-d, --debug           Print debug messages.
```

Exporting an IceStorm Database

To export an IceStorm database, use the `--export` option to specify the output file and the `--dbhome` option to specify the path name of the service's database environment. To discover the location of your database environment, review the service's configuration and look for the setting of `Freeze.DbEnv.IceStorm.DbHome`. For example, the IceStorm sample programs typically use this setting:

```
Freeze.DbEnv.IceStorm.DbHome=db
```

Run the following command to export the database:

```
$ icestormdb --export icestorm.ixp --dbhome db
```

The utility will fail with an error message if the given path does not contain an IceStorm database.

Importing an IceStorm Database

To import an IceStorm database, use the `--import` option to specify the input file and the `--dbhome` option to specify the path name of the service's database environment. For example, use the following command to import a database into the `dbNew` directory from a file named `icestorm.ixp`:

```
$ icestormdb --import icestorm.ixp --dbhome dbNew
```

The target directory must be empty.

Copyright 2016, ZeroC, Inc.

The Ice Protocol

The Ice protocol definition consists of three major parts:

- a set of data encoding rules that determine how the various data types are serialized
- a number of message types that are interchanged between client and server, together with rules as to what message is to be sent under what circumstances
- a set of rules that determine how client and server agree on a particular protocol and encoding version

Topics

- Data Encoding
- Protocol Messages
- Protocol Compression
- Protocol and Encoding Versions

Copyright 2016, ZeroC, Inc.

Data Encoding

The key goals of the Ice data encoding are simplicity and efficiency. In keeping with these principles, the encoding does not align primitive types on word boundaries and therefore eliminates the wasted space and additional complexity that alignment requires. The Ice data encoding simply produces a stream of contiguous bytes; data contains no padding bytes and need not be aligned on word boundaries.

Data is always encoded using little-endian byte order for numeric types. (Most machines use a little-endian byte order, so the Ice data encoding is "right" more often than not.) Ice does not use a "receiver makes it right" scheme because of the additional complexity this would introduce. Consider, for example, a chain of receivers that merely forward data along the chain until that data arrives at an ultimate receiver. (Such topologies are common for event distribution services.) The Ice protocol permits all the intermediates to forward the data without requiring it to be unmarshaled: the intermediates can forward requests by simply copying blocks of binary data. With a "receiver makes it right" scheme, the intermediates would have to unmarshal and remarshal the data whenever the byte order of the next receiver in the chain differs from the byte order of the sender, which is inefficient.

Ice requires clients and servers that run on big-endian machines to incur the extra cost of byte swapping data into little-endian layout, but that cost is insignificant compared to the overall cost of sending or receiving a request.

Topics

- Basic Data Encoding
- Data Encoding for Exceptions
- Data Encoding for Classes
- Data Encoding for Interfaces
- Data Encoding for Proxies
- Data Encoding for Optional Values

See Also

- Protocol Messages

Copyright 2016, ZeroC, Inc.

Basic Data Encoding

On this page:

- Encoding for Sizes
- Encoding for Encapsulations
- Encoding for Slices
- Encoding for Basic Types
- Encoding for Strings
- Encoding for Sequences
- Encoding for Dictionaries
- Encoding for Enumerators
- Encoding for Structures

Encoding for Sizes

Many of the types involved in the data encoding, as well as several protocol message components, have an associated size or count. A size is a non-negative number. Sizes and counts are encoded in one of two ways:

1. If the number of elements is less than 255, the size is encoded as a single `byte` indicating the number of elements.
2. If the number of elements is greater than or equal to 255, the size is encoded as a `byte` with value `255`, followed by an `int` indicating the number of elements.

Using this encoding to indicate sizes is significantly cheaper than always using an `int` to store the size, especially when marshaling sequences of short strings: counts of up to 254 require only a single byte instead of four. This comes at the expense of counts greater than 254, which require five bytes instead of four. However, for sequences or strings of length greater than 254, the extra byte is insignificant.

Encoding for Encapsulations

An encapsulation is used to contain variable-length data that an intermediate receiver may not be able to decode, but that the receiver can forward to another recipient for eventual decoding. An encapsulation is encoded as if it were the following structure:

Slice

```
struct Encapsulation {
    int size;
    byte major;
    byte minor;
    // [... size - 6 bytes ...]
};
```

The `size` member specifies the size of the encapsulation in bytes (including the `size`, `major`, and `minor` fields). The `major` and `minor` fields specify the encoding version of the data contained in the encapsulation. The version information is followed by `size-6` bytes of encoded data.

All the data in an encapsulation is context-free, that is, nothing inside an encapsulation can refer to anything outside the encapsulation. This property allows encapsulations to be forwarded among address spaces as a blob of data.

Encapsulations can be nested, that is, contain other encapsulations.

An encapsulation can be empty, in which case the value of `size` is 6.

Encoding for Slices

The encoding format of slices changed in version 1.1.

Slice Encoding version 1.0

Exceptions and classes may be subject to *slicing* if the receiver of a value only partially understands the received value (that is, only has

Copyright 2016, ZeroC, Inc.

knowledge of a base type, but not of the actual run-time derived type). To allow the receiver of an exception or class to ignore those parts of a value that it does not understand, exception and class values are marshaled as a sequence of slices (one slice for each level of the inheritance hierarchy). A slice is a byte count encoded as a fixed-length four-byte integer, followed by the data for the slice. (The byte count includes the four bytes occupied by the count itself, so an empty slice has a byte count of four and no data.) The receiver of a value can skip over a slice by reading the byte count *b*, and then discarding the next *b-4* bytes in the input stream.

Slice Encoding version 1.1

Version 1.1 of the encoding still marshals exceptions and classes as slices in conceptually the same manner as for version 1.0, but bit flags in the leading byte of each slice determine its format and content.
Type ID

The initial slice of a class or exception, representing the instance's most-derived type, always includes a type ID. For an exception, the type ID in the initial slice is encoded as a string. For a class, the type ID in the initial slice can either be encoded as a string, an index (if the same type ID has already been encoded in the current encapsulation), or a compact ID.

Whether any subsequent slices include some form of type ID depends on the format with which the value was encoded: to facilitate slicing an instance to a less-derived type, the sliced format includes a type ID in every slice, whereas the compact format excludes type IDs in subsequent slices to conserve space while sacrificing the slicing feature.
Optional Data Members

This flag is true if the slice includes any optional data members, which are encoded after all required members. If a slice encodes its size, the size includes the optional data members.
Object Indirection Table

This flag can only be true when using the sliced format. In this case, data members that refer to class instances are encoded as indices into an indirection table that immediately follows the slice. The slice's size does *not* include the indirection table. This flag should only be set to true when there is at least one non-nil object reference in the slice, that is, when the indirection table is not empty.
Slice Size

If this flag is true, a byte count encoded as a fixed-length four-byte integer immediately follows the type ID. (The byte count includes the four bytes occupied by the count itself, so an empty slice has a byte count of four and no data.) The byte count indicates the number of bytes occupied by the encoding for the required and optional data members, but does not include the size of an object indirection table, if present.

A receiver can skip over a slice by reading the byte count *b*, and then discarding the next *b-4* bytes in the input stream. If an object indirection table is present, the receiver must then decode the table.

If this flag is false, it implies that the sender used the compact format and therefore skipping slices is not possible. The receiver must know the most-derived type in this situation otherwise decoding will fail.
Last Slice

This flag indicates whether the current slice is the last slice of the instance.
Slice Flags

The table below shows how to interpret the bit flags in the leading byte of a slice:

Bit number	Description
0-1	0 = no type ID is encoded for the slice
	1 = type ID is encoded as a string
	2 = type ID is an index encoded as a size
	3 = type ID is a compact ID encoded as a size
2	Whether or not the slice contains optional data members
3	Whether or not the slice contains an indirection table
4	Whether or not the slice size follows the type ID
5	If 0, more slices will follow, if 1, this is the last slice
6	Reserved for future use
7	Reserved for future use

Bit flags for a slice.

Copyright 2016, ZeroC, Inc.

Encoding for Basic Types

The basic types are encoded as shown in the table. Integer types (`short`, `int`, `long`) are represented as two's complement numbers, and floating point types (`float`, `double`) use the IEEE standard formats [1]. All numeric types use a little-endian byte order.

Type	Encoding
`bool`	A single byte with value `1` for `true`, `0` for `false`
`byte`	An uninterpreted byte
`short`	Two bytes (LSB, MSB)
`int`	Four bytes (LSB .. MSB)
`long`	Eight bytes (LSB .. MSB)
`float`	Four bytes (23-bit fractional mantissa, 8-bit exponent, sign bit)
`double`	Eight bytes (52-bit fractional mantissa, 11-bit exponent, sign bit)

Encoding for basic types.

Encoding for Strings

Strings are encoded as a size, followed by the string contents in UTF-8 format [2]. Strings are not NUL-terminated. An empty string is encoded with a size of zero.

Encoding for Sequences

Sequences are encoded as a size representing the number of elements in the sequence, followed by the elements encoded as specified for their type.

Encoding for Dictionaries

Dictionaries are encoded as a size representing the number of key-value pairs in the dictionary, followed by the pairs. Each key-value pair is encoded as if it were a `struct` containing the key and value as members, in that order.

Encoding for Enumerators

The encoding format of enumerators changed in version 1.1.

Enumerator Encoding version 1.0

The number of bytes required to encode an enumerator in version 1.0 is determined by the largest value in the enumeration. In enumerations with no custom enumerator values, the largest value is the number of enumerators less one; the value encoded for an enumerator is its ordinal value in the definition, with the first enumerator having the value zero. For example, the largest value in the following enumeration is 2:

Slice

```
enum Fruit { Apple, Pear, Orange }; // Encoded values: Apple = 0, Pear = 1, Orange
= 2
```

When an enumeration includes custom values, the encoding uses the enumerator's assigned value, which is not necessarily the same as its ordinal value. Consider this example:

Copyright 2016, ZeroC, Inc.

Slice

```
enum Fruit { Apple = 1, Pear = 3, Orange }; // Encoded values: Apple = 1, Pear = 3,
Orange = 4
```

The largest value in this case is 4.

An enumerator is encoded as follows:

- If the largest value is in the range 0..126, the enumerator's value is marshaled as a `byte`.
- If the largest value is in the range 127..32766, the enumerator's value is marshaled as a `short`.
- If the largest value is greater than 32766, the enumerator's value is marshaled as an `int`.

Changing the definition of an enumeration can break compatibility with existing applications. For example, if enumerators are added, removed, or changed such that the largest value crosses one of the thresholds shown above, the encoded form of the enumerators will change and cause marshaling errors unless you rebuild all applications that use this definition.

Enumerator Encoding version 1.1

An enumerator is encoded as a size, meaning the encoding of an enumerator requires one byte if its value is less than 255, or five bytes if its value is 255 or greater. The encoding uses the Slice value of the enumerator, which is not necessarily the same as its ordinal value. Repeating the previous example:

Slice

```
enum Fruit { Apple = 1, Pear = 3, Orange }; // Encoded values: Apple = 1, Pear = 3,
Orange = 4
```

Although enumerator `Pear` may have ordinal value 1 in some language mappings (notably, Java), the encoding uses its Slice value of 3.

Encoding for Structures

The members of a structure are encoded in the order they appear in the `struct` declaration, as specified for their types.

See Also

- Protocol Messages
- Protocol and Encoding Versions
- Data Encoding for Exceptions
- Data Encoding for Classes

References

1. Institute of Electrical and Electronics Engineers. 1985. *IEEE 754-1985 Standard for Binary Floating-Point Arithmetic*. Piscataway, NJ: Institute of Electrical and Electronic Engineers.
2. Unicode Consortium, ed. 2000. *The Unicode Standard, Version 3.0*. Reading, MA: Addison-Wesley.

Copyright 2016, ZeroC, Inc.

Data Encoding for Exceptions

Exception Encoding version 1.0

An exception is marshaled as shown below:

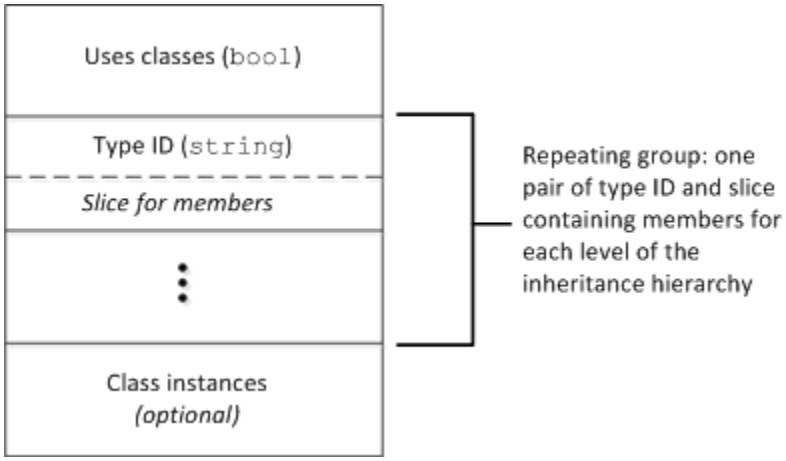

Marshaling format for exceptions.

Every exception instance is preceded by a single byte that indicates whether the exception uses class members: the byte value is `1` if any of the exception members are classes (or if any of the exception members, recursively, contain class members) and `0`, otherwise.

Following the header byte, the exception is marshaled as a sequence of pairs: the first member of each pair is the type ID for an exception slice, and the second member of the pair is a slice containing the marshaled members of that slice. The sequence of pairs is marshaled in derived-to-base order, with the most-derived slice first, and ending with the least-derived slice. Within each slice, data members are marshaled as for structures: in the order in which they are defined in the Slice definition.

Following the sequence of pairs, any class instances that are used by the members of the exception are marshaled. This final part is optional: it is present only if the header byte is `1`.

To illustrate the marshaling, consider the following exception hierarchy:

Slice

```
exception Base {
    int baseInt;
    string baseString;
};

exception Derived extends Base {
    bool derivedBool;
    string derivedString;
    double derivedDouble;
};
```

Assume that the exception members are initialized to the values shown below:

Member	Type	Value	Marshaled size (in bytes)
`baseInt`	`int`	`99`	`4`
`baseString`	`string`	`"Hello"`	`6`
`derivedBool`	`bool`	`true`	`1`

Copyright 2016, ZeroC, Inc.

`derivedString`	`string`	`"World!"`	7
`derivedDouble`	`double`	`3.14`	8

Member values of an exception of type `Derived`.

From the above table, we can see that the total size of the members of `Base` is 10 bytes, and the total size of the members of `Derived` is 16 bytes. None of the exception members are classes. An instance of this exception has the on-the-wire representation shown in the next table. (The size, type, and byte offset of the marshaled representation is indicated for each component.)

Marshaled value	Size in bytes	Type	Byte offset
`0` *(no class members)*	1	`bool`	0
`"::Derived"` *(type ID)*	10	`string`	1
`20` *(byte count for slice)*	4	`int`	11
`1` (*`derivedBool`*)	1	`bool`	15
`"World!"` (*`derivedString`*)	7	`string`	16
`3.14` (*`derivedDouble`*)	8	`double`	23
`"::Base"` *(type ID)*	7	`string`	31
`14` *(byte count for slice)*	4	`int`	38
`99` (*`baseInt`*)	4	`int`	42
`"Hello"` (*`baseString`*)	6	`string`	46

Marshaled representation of the exception.

Note that the size of each string is one larger than the actual string length. This is because each string is preceded by a count of its number of bytes, as directed by the encoding for strings.

The receiver of this sequence of values uses the header byte to decide whether it eventually must unmarshal any class instances contained in the exception (none in this example) and then examines the first type ID (`::Derived`). If the receiver recognizes that type ID, it can unmarshal the contents of the first slice, followed by the remaining slices; otherwise, the receiver reads the byte count that follows the unknown type (20) and then skips 20-4 bytes in the input stream, which is the start of the type ID for the second slice (`::Base`). If the receiver does not recognize that type ID either, it again reads the byte count following the type ID (14), skips 14-4 bytes, and attempts to read another type ID. (This can happen only if client and server have been compiled with mismatched Slice definitions that disagree in the exception specification of an operation.) In this case, the receiver will eventually encounter an unmarshaling error, which it can report with a `MarshalException`.

If an exception contains class members, these members are marshaled following the exception slices as described in the class encoding.

Exception Encoding version 1.1

An exception is marshaled as a collection of slices whose order matches the inheritance hierarchy, with the most-derived type appearing first. The selected encoding format affects the content of each slice. The final slice, representing the least-derived type, has its *last slice* bit set to true.

An exception in the compact format is marshaled as follows:

Copyright 2016, ZeroC, Inc.

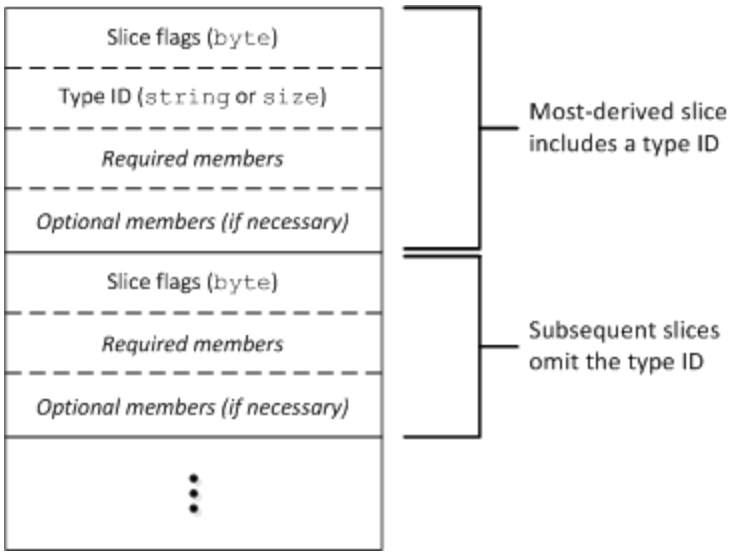

Compact format for exceptions.

The leading byte of each slice is a set of bit flags that specifies the features of the slice. The compact format includes a type ID in the initial (most-derived) slice but omits the type ID from all subsequent slices.

The sliced format includes a type ID in every slice, along with a slice size and an optional indirection table:

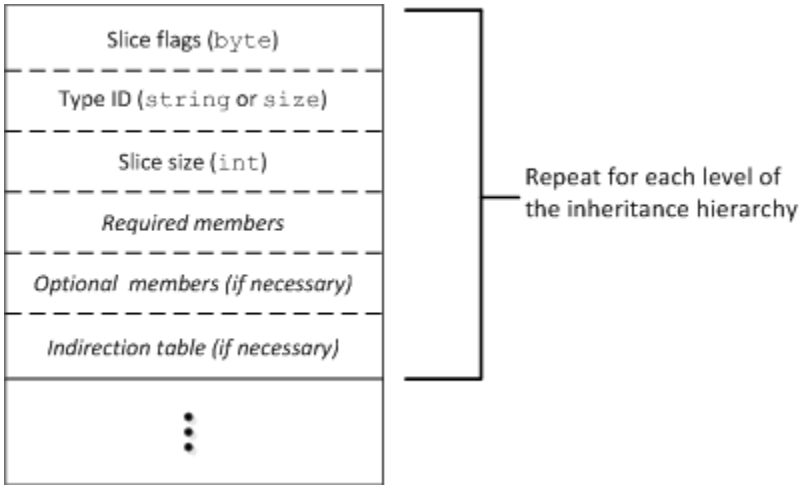

Sliced format for exceptions.

To illustrate the marshaling, consider the following exception hierarchy:

Copyright 2016, ZeroC, Inc.

Slice

```
exception Base {
    int baseInt;
    string baseString;
};

exception Derived extends Base {
    bool derivedBool;
    string derivedString;
    double derivedDouble;
};
```

Assume that the exception members are initialized to the values shown below:

Member	Type	Value	Marshaled size (in bytes)
`baseInt`	`int`	`99`	4
`baseString`	`string`	`"Hello"`	6
`derivedBool`	`bool`	`true`	1
`derivedString`	`string`	`"World!"`	7
`derivedDouble`	`double`	`3.14`	8

Member values of an exception of type `Derived`.

From the above table, we can see that the total size of the members of `Base` is 10 bytes, and the total size of the members of `Derived` is 16 bytes. None of the exception members are classes. An instance of this exception using the sliced format has the on-the-wire representation shown in the next table. (The size, type, and byte offset of the marshaled representation is indicated for each component.)

Marshaled value	Size in bytes	Type	Byte offset
`18` *(flags: type ID is a string, slice size is present)*	1	`byte`	0
`"::Derived"` *(type ID)*	10	`string`	1
`20` *(byte count for slice)*	4	`int`	11
`1` *(`derivedBool`)*	1	`bool`	15
`"World!"` *(`derivedString`)*	7	`string`	16
`3.14` *(`derivedDouble`)*	8	`double`	23
`50` *(flags: type ID is a string, slice size is present, last slice)*	1	`byte`	31
`"::Base"` *(type ID)*	7	`string`	32
`14` *(byte count for slice)*	4	`int`	39
`99` *(`baseInt`)*	4	`int`	43
`"Hello"` *(`baseString`)*	6	`string`	47

Marshaled representation of the exception using the sliced format.

Note that the size of each string is one larger than the actual string length. This is because each string is preceded by a count of its number of bytes, as directed by the encoding for strings.

Repeating this exercise using the compact format produces the following encoding:

Marshaled value	Size in bytes	Type	Byte offset

Copyright 2016, ZeroC, Inc.

2 *(flags: type ID is a string)*	1	`byte`	0
`"::Derived"` *(type ID)*	10	`string`	1
1 *(`derivedBool`)*	1	`bool`	11
`"World!"` *(`derivedString`)*	7	`string`	12
3.14 *(`derivedDouble`)*	8	`double`	19
32 *(flags: last slice)*	1	`byte`	27
99 *(`baseInt`)*	4	`int`	28
`"Hello"` *(`baseString`)*	6	`string`	32

Marshaled representation of the exception using the compact format.

When using the compact format, the receiver *must* know the most-derived type: the only type ID included in the encoding is that of the most-derived type. Furthermore, the lack of slice sizes means the receiver cannot skip a slice without knowing how to decode its contents.

See Also

- Type IDs
- Basic Data Encoding
- Data Encoding for Classes

Copyright 2016, ZeroC, Inc.

Data Encoding for Classes

The marshaling for classes is complex, due to the need to deal with the pointer semantics for graphs of classes, as well as the need for the receiver to slice or preserve classes of unknown derived type. In addition, the marshaling for classes uses a type ID compression scheme to avoid repeatedly marshaling the same type IDs for large graphs of class instances.

Encoding for Class References

Consider the following Slice definitions:

Slice

```
class Node {
    int value;
    Node next;
};

struct S {
    Node obj;
};
```

We call the `obj` member a *reference* to a class instance. References can appear as data members, as in the case of `obj` above, or nested inside of other types, such as a sequence element or dictionary value. There are significant differences in the encoding for references between versions 1.0 and 1.1.

Class Reference Encoding version 1.0

Ice encodes a nil reference as a 32-bit integer with value zero. For non-nil references, the encoder maintains a table per encapsulation that associates a unique non-zero positive integer with each class instance. We refer to this integer as an *instance ID*. The first time the encoder encounters a reference to a particular instance, it assigns the instance an unused ID, inserts it into the table, and encodes the reference as a 32-bit integer whose value is the *negative* of the ID. If the encoder has already encountered that instance, it simply encodes the instance's previously-assigned ID as a negative 32-bit integer.

All class instances are encoded at the end of the encapsulation.

Class Reference Encoding version 1.1

As in version 1.0 of the encoding, each instance is assigned a unique ID per encapsulation. However, to conserve space, version 1.1 encodes a reference as a size value, with a nil reference encoded as a size value of zero and a non-nil reference encoded as a positive size value. The encoding reserves ID value zero to denote a nil reference, and ID value 1 to denote an *inline* instance in which an instance's state is marshaled at the point of its first reference. Consequently, instance IDs must start at 2. The encoding assigns instance IDs sequentially in order of appearance, with the first instance assigned an ID of 2, the next instance has ID 3, and so on.

The encoding for a reference depends on the format being used, and may also depend on the context in which the reference occurs:

- **Compact format**
  In this format, the encoding always uses the inline scheme. Suppose we are encoding a value of the structure type `S` shown earlier. If its `obj` member refers to a class instance that the encoder has not yet encountered in the current encapsulation, the encoding assigns it the next available sequential ID, marshals a size with value 1 signifying that an instance is about to be written, and then immediately marshals the instance itself. Any subsequent reference to the same instance is encoded as a size whose value is the instance's corresponding ID.
- **Sliced format**
  When a reference occurs *outside* the context of a class instance, Ice uses the same inline scheme as in the compact format. For example, suppose a Slice operation declared a parameter of structure type `S` shown above. When marshaling this parameter, its `obj` member is encoded as a size with value 1, followed by the encoding of the `Node` instance itself. Now consider the marshaling of the `Node` instance, where we encounter a class reference in member `next` that occurs *inside* the context of a class instance. To assist the Ice run time in unmarshaling and remarshaling instances, such a reference is encoded as a size whose value is an index into an indirection table.

Copyright 2016, ZeroC, Inc.

Encoding for Class Instances

It is important to understand the distinction between marshaling a reference and marshaling an instance. In C++ terms, this is equivalent to the difference between a pointer and the heap data containing an object's state at which the pointer is pointing. Marshaling an instance means we are encoding the data members of a class instance.

Class Instance Encoding version 1.0

Classes are marshaled as a number of pairs containing a type ID and a slice (one pair for each level of the inheritance hierarchy) and marshaled in derived-to-base order. Only data members are marshaled — no information is sent that would relate to operations. Each marshaled class instance is preceded by a (non-zero) positive integer that provides an identity for the instance. The sender assigns this identity during marshaling such that each marshaled instance has a different identity. The receiver uses that identity to correctly reconstruct graphs of classes. The overall marshaling format for classes is shown below:

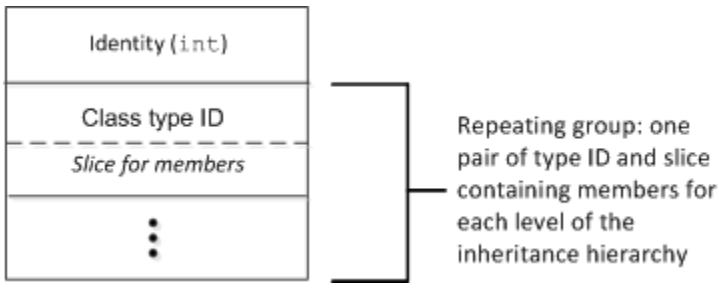

Marshaling format for classes.

Prior to ending an encapsulation, the sender must encode all instances referenced within that encapsulation in an object table. The encoding may require multiple passes: as the sender processes an instance in the current pass, it may encounter additional instances to be processed in a subsequent pass. For each pass, the sender encodes a size denoting the number of objects in the pass. To signify that all instances have been encoded, the sender must marshal a size value of zero.

This object table must be written if any operation parameter or data member has a class type, even if all references are nil.

Class Instance Encoding version 1.1

The leading byte of a class instance is a size value of 1. Following this byte is a collection of slices arranged in derived-to-base order. Only data members are marshaled — no information is sent that would relate to operations. The initial (most-derived) slice always includes a type ID that may be encoded as a string or as a numeric value, as specified by the flags that begin each slice. Depending on the format being used, subsequent slices may or may not include a type ID.

Each slice consists of a leading byte representing the slice flags, an optional type ID, an optional slice size, the required members for that slice in order of declaration, and the optional members of that slice. The sender must set the appropriate slice flags to indicate whether the slice includes a type ID, size, and optional data members, and whether this is the last slice of the instance. The compact format only includes a type ID in the initial (most-derived) slice and omits the slice size, as shown in the following diagram:

Copyright 2016, ZeroC, Inc.

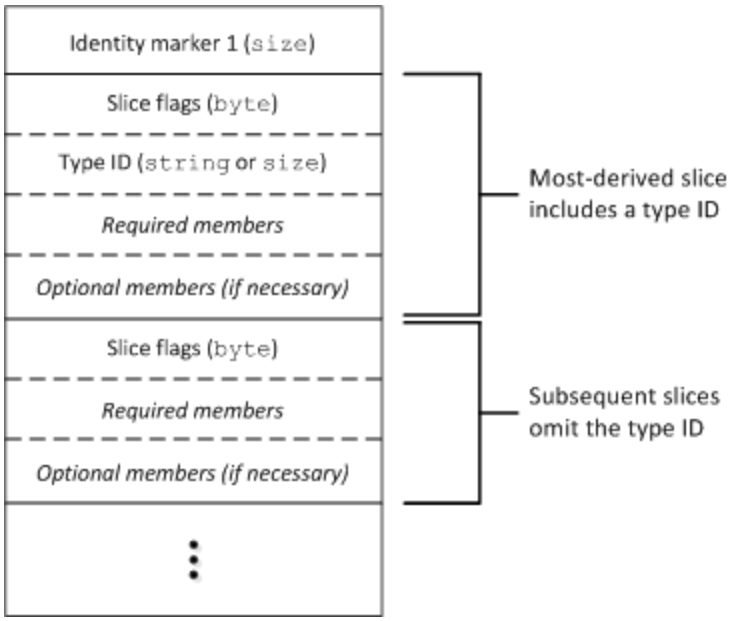

Compact format for classes.

When using the sliced format, a type ID is included in every slice, along with a slice size:

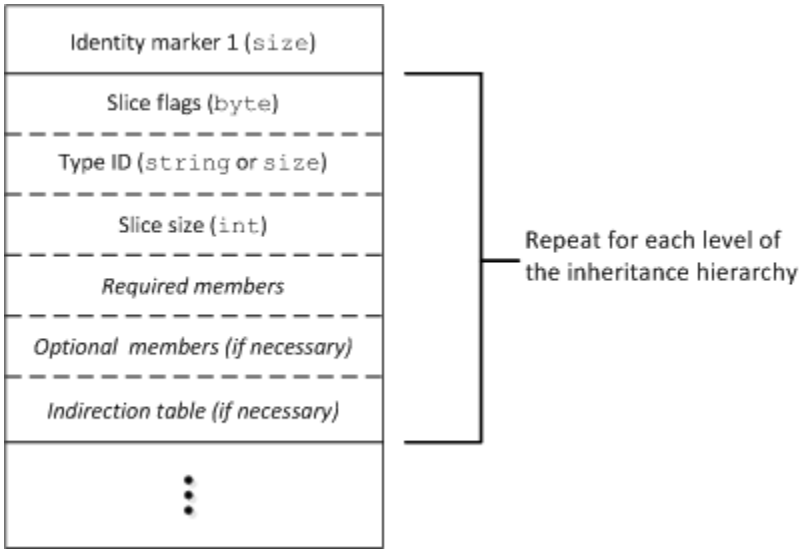

Sliced format for classes.

A slice that contains at least one non-nil class reference must use an indirection table for all references in that slice. The table is an array of instances encoded using the inline scheme. A leading size value indicates the number of elements in the table; each element is either an instance ID (if the instance has already been encoded within the current encapsulation), or the instance itself denoted by a leading size of 1. The indirection table appears in the encoding immediately following the required and optional data members, *but is not included in the byte count denoted by the slice size*. To skip a slice for an unknown type, the receiver can advance the input stream by the number of bytes specified in the slice size, but then must process all references or instances in the indirection table, if one is present.

To support slice preservation for an instance, the receiver must temporarily retain the slices of any unknown derived types, and also be able to reconstruct the indirection table in its original order for each of these slices in case the instance is later remarshaled.

See Also

- Classes
- Type IDs
- Data Encoding for Exceptions

Copyright 2016, ZeroC, Inc.

- Basic Data Encoding
- Slicing Values and Exceptions

Copyright 2016, ZeroC, Inc.

Data Encoding for Class Type IDs

Type ID Encoding version 1.0

Unlike for exception type IDs, class type IDs are not simple strings. Instead, a class type ID is marshaled as a boolean followed by either a string or a size, to conserve bandwidth. To illustrate this, consider the following class hierarchy:

Slice

```
class Base {
    // ...
};

class Derived extends Base {
    // ...
};
```

The type IDs for the class slices are `::Derived` and `::Base`. Suppose the sender marshals three instances of `::Derived` as part of a single request. (For example, two instances could be out-parameters and one instance could be the return value.)

The first instance that is sent on the wire contains the type IDs `::Derived` and `::Base` preceding their respective slices. Because marshaling proceeds in derived-to-base order, the first type ID that is sent is `::Derived`. Every time the sender sends a type ID that it has not sent previously in the same request, it sends the boolean value `false`, followed by the type ID. Internally, the sender also assigns a unique positive number to each type ID. These numbers start at `1` and increment by one for each type ID that has not been marshaled previously. This means that the first type ID is encoded as the boolean value `false`, followed by `::Derived`, and the second type ID is encoded as the boolean value `false`, followed by `::Base`.

When the sender marshals the remaining two instances, it consults a lookup table of previously-marshaled type IDs. Because both type IDs were sent previously in the same request (or reply), the sender encodes all further occurrences of `::Derived` as the value `true` followed by the number `1` encoded as a size, and it encodes all further occurrences of `::Base` as the value `true` followed by the number `2` encoded as a size.

When the receiver reads a type ID, it first reads its boolean marker:

- If the boolean is `false`, the receiver reads a string and enters that string into a lookup table that maps integers to strings. The first new class type ID received in a request is numbered `1`, the second new class type ID is numbered `2`, and so on.
- If the boolean value is `true`, the receiver reads a number encoded as a size and uses it to retrieve the corresponding class type ID from the lookup table.

Note that this numbering scheme is re-established for each new encapsulation. (As we will see in our discussion of protocol messages, parameters, return values, and exceptions are always marshaled inside an enclosing encapsulation.) For subsequent or nested encapsulation, the numbering scheme restarts, with the first new type ID being assigned the value `1`. In other words, each encapsulation uses its own independent numbering scheme for class type IDs to satisfy the constraint that encapsulations must not depend on their surrounding context.

Encoding class type IDs in this way provides significant savings in bandwidth: whenever an ID is marshaled a second and subsequent time, it is marshaled as a two-byte value (assuming no more than 254 distinct type IDs per request) instead of as a string. Because type IDs can be long, especially if you are using nested modules, the savings are considerable.

Type ID Encoding version 1.1

Each slice of a class instance has a leading byte containing flags that describe various aspects of the slice, including whether the slice includes a type ID and how that type ID is encoded. There are four possibilities:

1. No type ID included
2. Type ID is encoded as a string
3. Type ID is encoded as an index
4. Type ID is encoded as a compact ID

The initial slice of an instance, representing the most-derived type, always contains some form of type ID so that the receiver knows what Slice type is present. Depending on the format used by the sender, subsequent slices may or may not include a type ID: the compact format omits type IDs in subsequent slices, whereas the sliced format includes a type ID in every slice. A receiver need only examine the slice flags

Copyright 2016, ZeroC, Inc.

to discover how to decode the type ID.

String Type IDs

The encoding for string type IDs uses a "compression" scheme similar to that of version 1.0: within an encapsulation, a given type ID is never encoded as a string more than once. The first time a sender encounters a type ID, the sender assigns an integer index to the ID and encodes the ID as a string. For all subsequent occurrences of the same type ID within the encapsulation, the sender encodes the index associated with that type ID as a size. Index values start at 1 and increase sequentially with each new type ID. The sender is responsible for setting the relevant bits in the flags of each slice to specify how the type ID is encoded.

The slice flags in version 1.1 of the encoding serve the same purpose as the boolean value that precedes each type ID in version 1.0, without consuming an entire byte.

Compact Type IDs

As a way of further reducing the overhead associated with class instances, version 1.1 adds the ability to substitute numeric values for strings when encoding type IDs. Consider the following Slice definitions:

Slice

```
class Base(3) {
    // ...
};

class Derived(4) extends Base {
    // ...
};

class MoreDerived extends Derived {
    // ...
};
```

The value in parentheses after the class name represents the *compact* type ID for the class. The sender's format determines whether each slice includes a type ID. If a given slice includes a type ID, the encoding always uses a compact type ID (if defined) in preference to its string equivalent. Suppose a sender is encoding an instance of `MoreDerived` in the sliced format. The initial slice uses the string type ID for `MoreDerived`, the next slice uses the compact type ID for `Derived`, and the last slice uses the compact type ID for `Base`.

A compact type ID is encoded as a size, with the relevant bits set in the slice flags.

The developer is responsible for ensuring compact type IDs are sufficiently unique. Using values less than 255 produces the most efficient encoding.

See Also

- Type IDs
- Basic Data Encoding
- Protocol Messages

Copyright 2016, ZeroC, Inc.

Simple Example of Class Encoding

On this page:

- Sample Class Definitions
- Class Encoding version 1.0
- Class Encoding version 1.1
  - Class Encoding in the Sliced Format
  - Class Encoding in the Compact Format
  - Class Encoding in the Compact Format with Compact Type IDs

Sample Class Definitions

We have separately discussed the primary components of the class encoding: slices, references, and type IDs. To make the preceding discussions more concrete, consider the following class definitions:

Slice

```
interface SomeInterface {
     void op1();
};

class Base {
    int baseInt;
    void op2();
    string baseString;
};

class Derived extends Base implements SomeInterface {
    bool derivedBool;
    string derivedString;
    void op3();
    double derivedDouble;
};
```

Note that `Base` and `Derived` have operations, and that `Derived` also implements the interface `SomeInterface`. Because marshaling of classes is concerned with state, not behavior, the operations `op1`, `op2`, and `op3` are simply ignored during marshaling and the on-the-wire representation is as if the classes had been defined as follows:

Slice

```
class Base {
    int baseInt;
    string baseString;
};

class Derived extends Base {
    bool derivedBool;
    string derivedString;
    double derivedDouble;
};
```

Suppose the sender marshals two instances of `Derived` (for example, as two in-parameters in the same request) with these member

Copyright 2016, ZeroC, Inc.

values:

First instance:

Member	Type	Value	Marshaled size (in bytes)
`baseInt`	`int`	`99`	4
`baseString`	`string`	`"Hello"`	6
`derivedBool`	`bool`	`true`	1
`derivedString`	`string`	`"World!"`	7
`derivedDouble`	`double`	`3.14`	8

Second instance:

Member	Type	Value	Marshaled size (in bytes)
`baseInt`	`int`	`115`	4
`baseString`	`string`	`"Cave"`	5
`derivedBool`	`bool`	`false`	1
`derivedString`	`string`	`"Canem"`	6
`derivedDouble`	`double`	`6.32`	8

We describe how to marshal these instances using versions 1.0 and 1.1 of the encoding in separate sections below.

Class Encoding version 1.0

The sender arbitrarily assigns a non-zero identity to each instance. Typically, the sender will simply consecutively number the instances starting at 1. For this example, assume that the two instances have the identities 1 and 2. The marshaled representation for the two instances (assuming that they are marshaled immediately following each other) is shown below:

Marshaled value	Size in bytes	Type	Byte offset
`1` *(identity)*	4	`int`	0
`0` *(marker for class type ID)*	1	`bool`	4
`"::Derived"` *(class type ID)*	10	`string`	5
`20` *(byte count for slice)*	4	`int`	15
`1` (*`derivedBool`*)	1	`bool`	19
`"World!"` (*`derivedString`*)	7	`string`	20
`3.14` (*`derivedDouble`*)	8	`double`	27
`0` *(marker for class type ID)*	1	`bool`	35
`"::Base"` *(type ID)*	7	`string`	36
`14` *(byte count for slice)*	4	`int`	43
`99` (*`baseInt`*)	4	`int`	47
`"Hello"` (*`baseString`*)	6	`string`	51
`0` *(marker for class type ID)*	1	`bool`	57
`"::Ice::Object"` *(class type ID)*	14	`string`	58
`5` *(byte count for slice)*	4	`int`	72
`0` *(number of dictionary entries)*	1	`size`	76
`2` *(identity)*	4	`int`	77
`1` *(marker for class type ID)*	1	`bool`	81

Copyright 2016, ZeroC, Inc.

`1` *(class type ID)*	`1`	`size`	82
`19` *(byte count for slice)*	`4`	`int`	83
`0` *(`derivedBool`)*	`1`	`bool`	87
`"Canem"` *(`derivedString`)*	`6`	`string`	88
`6.32` *(`derivedDouble`)*	`8`	`double`	94
`1` *(marker for class type ID)*	`1`	`bool`	102
`2` *(class type ID)*	`1`	`size`	103
`13` *(byte count for slice)*	`4`	`int`	104
`115` *(`baseInt`)*	`4`	`int`	108
`"Cave"` *(`baseString`)*	`5`	`string`	112
`1` *(marker for class type ID)*	`1`	`bool`	117
`3` *(class type ID)*	`1`	`size`	118
`5` *(byte count for slice)*	`4`	`int`	119
`0` *(number of dictionary entries)*	`1`	`size`	123

Note that, because classes (like exceptions) are sent as a sequence of slices, the receiver of a class can slice off any derived parts of a class it does not understand. Also note that (as shown in the above table) each class instance contains three slices. The third slice is for the type `::Ice::Object`, which is the base type of all classes. The class type ID `::Ice::Object` has the number 3 in this example because it is the third distinct type ID that is marshaled by the sender. (See entries at byte offsets 58 and 118 in the above table.) All class instances have this final slice of type `::Ice::Object`.

Marshaling a separate slice for `::Ice::Object` dates back to Ice versions 1.3 and earlier. In those versions, classes carried a facet map that was marshaled as if it were defined as follows:

Slice

```
module Ice {
    class Object;

    dictionary<string, Object> FacetMap;

    class Object {
        FacetMap facets; // No longer exists
    };
};
```

As of Ice version 1.4, this facet map is always empty, that is, the count of entries for the dictionary that is marshaled in the `::Ice::Object` slice is always zero. If a receiver receives a class instance with a non-empty facet map, it must throw a `MarshalException`.

Note that if a class has no data members, a type ID and slice for that class is still marshaled. The byte count of the slice will be 4 in this case, indicating that the slice contains no data.

Class Encoding version 1.1

A leading size value of `1` marks the beginning of an instance, followed by one or more slices.

Class Encoding in the Sliced Format

The marshaled representation for the two instances (assuming that they are marshaled immediately following each other) in the sliced format is shown below:

Copyright 2016, ZeroC, Inc.

Marshaled value	Size in bytes	Type	Byte offset
`1` *(instance marker)*	1	`size`	0
`17` *(slice flags: string type ID, size is present)*	1	`byte`	1
`"::Derived"` *(type ID - assigned index 1)*	10	`string`	2
`20` *(byte count for slice)*	4	`int`	12
`1` (*`derivedBool`*)	1	`bool`	16
`"World!"` (*`derivedString`*)	7	`string`	17
`3.14` (*`derivedDouble`*)	8	`double`	24
`49` *(slice flags: string type ID, size is present, last slice)*	1	`byte`	32
`"::Base"` *(type ID - assigned index 2)*	7	`string`	33
`14` *(byte count for slice)*	4	`int`	40
`99` (*`baseInt`*)	4	`int`	44
`"Hello"` (*`baseString`*)	6	`string`	48
`1` *(instance marker)*	1	`size`	54
`18` *(slice flags: index type ID, size is present)*	1	`byte`	55
`1` *(type ID index for `Derived`)*	1	`size`	56
`19` *(byte count for slice)*	4	`int`	57
`0` (*`derivedBool`*)	1	`bool`	61
`"Canem"` (*`derivedString`*)	6	`string`	62
`6.32` (*`derivedDouble`*)	8	`double`	68
`50` *(slice flags: index type ID, size is present, last slice)*	1	`byte`	76
`2` *(type ID index for `Base`)*	1	`size`	77
`13` *(byte count for slice)*	4	`int`	78
`115` (*`baseInt`*)	4	`int`	82
`"Cave"` (*`baseString`*)	5	`string`	86

The sliced format allows the receiver of a class to slice off any derived parts of a class it does not understand, as in version 1.0 of the encoding. Although the sliced format provides equivalent functionality to that of version 1.0, it is significantly more efficient, requiring only 91 bytes to encode our example compared to the 124 bytes required by version 1.0. We could reduce the encoded size even further, while still retaining the ability to slice off unknown types, by using compact type IDs.

Note that if a class has no data members, a type ID and slice for that class is still marshaled. The byte count of the slice will be 4 in this case, indicating that the slice contains no data.

Class Encoding in the Compact Format

The marshaled representation for the two instances (assuming that they are marshaled immediately following each other) in the compact format is shown below:

Marshaled value	Size in bytes	Type	Byte offset
`1` *(instance marker)*	1	`size`	0
`1` *(slice flags: string type ID)*	1	`byte`	1
`"::Derived"` *(type ID - assigned index 1)*	10	`string`	2
`1` (*`derivedBool`*)	1	`bool`	12
`"World!"` (*`derivedString`*)	7	`string`	13
`3.14` (*`derivedDouble`*)	8	`double`	20

Copyright 2016, ZeroC, Inc.

32 *(slice flags: last slice)*	1	byte	28
99 *(baseInt)*	4	int	29
"Hello" *(baseString)*	6	string	33
1 *(instance marker)*	1	size	39
2 *(slice flags: index type ID)*	1	byte	40
1 *(type ID index for Derived)*	1	size	41
0 *(derivedBool)*	1	bool	42
"Canem" *(derivedString)*	6	string	43
6.32 *(derivedDouble)*	8	double	49
32 *(slice flags: last slice)*	1	byte	57
115 *(baseInt)*	4	int	58
"Cave" *(baseString)*	5	string	62

In an effort to conserve bandwidth, the compact format omits certain details that would allow a receiver to slice off derived parts of a class, such as the slice size and the type IDs for base classes. The result is an encoding that requires only 67 bytes for the two sample instances.

Note that if a class has no data members, a type ID and slice for that class is still marshaled. The byte count of the slice will be 4 in this case, indicating that the slice contains no data.

Class Encoding in the Compact Format with Compact Type IDs

Compact type IDs can be used regardless of the sender's chosen format. For the sake of example, we will use compact type IDs together with the compact format to produce the smallest encoding possible. The Slice definitions below reflect the addition of the compact type IDs:

Slice

```
interface SomeInterface {
     void op1();
};

class Base(10) {
    int baseInt;
    void op2();
    string baseString;
};

class Derived(11) extends Base implements SomeInterface {
    bool derivedBool;
    string derivedString;
    void op3();
    double derivedDouble;
};
```

We assign the compact type ID `10` to `Base` and `11` to `Derived`. Note however that assigning a compact type ID to `Base` does not affect the size of the encoded data in our example because the compact format omits type IDs altogether for base types.

The marshaled representation for the two instances (assuming that they are marshaled immediately following each other) in the compact format is shown below:

Marshaled value	Size in bytes	Type	Byte offset
1 *(instance marker)*	1	size	0

Copyright 2016, ZeroC, Inc.

`3` *(slice flags: compact type ID)*	`1`	`byte`	`1`
`11` *(compact type ID for `Derived`)*	`1`	`size`	`2`
`1` *(`derivedBool`)*	`1`	`bool`	`3`
`"World!"` *(`derivedString`)*	`7`	`string`	`4`
`3.14` *(`derivedDouble`)*	`8`	`double`	`11`
`32` *(slice flags: last slice)*	`1`	`byte`	`19`
`99` *(`baseInt`)*	`4`	`int`	`20`
`"Hello"` *(`baseString`)*	`6`	`string`	`24`
`1` *(instance marker)*	`1`	`size`	`30`
`3` *(slice flags: compact type ID)*	`1`	`byte`	`31`
`11` *(compact type ID for `Derived`)*	`1`	`size`	`32`
`0` *(`derivedBool`)*	`1`	`bool`	`33`
`"Canem"` *(`derivedString`)*	`6`	`string`	`34`
`6.32` *(`derivedDouble`)*	`8`	`double`	`40`
`32` *(slice flags: last slice)*	`1`	`byte`	`48`
`115` *(`baseInt`)*	`4`	`int`	`49`
`"Cave"` *(`baseString`)*	`5`	`string`	`53`

Substituting a compact type ID for its string equivalent reduces the encoded size for the two instances by another nine bytes to 58, less than half the size of version 1.0.

See Also

- Data Encoding for Classes
- Data Encoding for Exceptions
- Basic Data Encoding
- Type IDs

Copyright 2016, ZeroC, Inc.

Data Encoding for Class Graphs

Classes support pointer semantics, that is, you can construct graphs of classes. It follows that classes can arbitrarily point at each other, and therefore the encoding must provide a scheme for serializing and deserializing a class graph. Note that the marshaling format for class references and instances differs significantly between versions 1.0 and 1.1 of the encoding.

On this page:

- Encoding a Class Graph version 1.0
- Encoding a Class Graph version 1.1
  - Class Graph in the Compact Format
  - Class Graph in the Sliced Format
  - Importance of the Indirection Table
- Impact of Slicing on Class Graph Decoding

Encoding a Class Graph version 1.0

In version 1.0 of the encoding, an instance ID is used to distinguish instances and pointers as follows:

- An instance ID of 0 denotes a null pointer.
- An instance ID > 0 precedes the marshaled contents of an instance.
- An instance ID < 0 denotes a pointer to an instance.

Instance ID values less than zero are pointers. For example, if the receiver receives the instance ID `-57`, this means that the corresponding class member that is currently being unmarshaled will eventually point at the instance with ID `57`.

For structures, classes, exceptions, sequences, and dictionary members that do not contain class members, the Ice encoding uses a simple depth-first traversal algorithm to marshal the members. For example, structure members are marshaled in the order of their Slice definition; if a structure member itself is of complex type, such as a sequence, the sequence is marshaled in toto where it appears inside its enclosing structure. For complex types that contain class members, this depth-first marshaling is suspended: instead of marshaling the actual class instance at this point, a negative instance ID is marshaled that indicates which class instance that member must eventually denote. For example, consider the following definitions:

Slice

```
class C {
    // ...
};

struct S {
    int i;
    C firstC;
    C secondC;
    C thirdC;
    int j;
};
```

Suppose we initialize a structure of type `S` as follows:

Copyright 2016, ZeroC, Inc.

C++

```
S myS;
myS.i = 99;
myS.firstC = new C;             // New instance
myS.secondC = 0;                // null
myS.thirdC = myS.firstC;        // Same instance as previously
myS.j = 100;
```

When this structure is marshaled, the contents of the three class members are not marshaled in-line. Instead, the sender marshals the negative instance IDs of the corresponding instances. Assuming that the sender has assigned the instance ID `78` to the instance assigned to `myS.firstC`, `myS` is marshaled as shown in the table.

Marshaled Value	Size in Bytes	Type	Byte offset
`99` *(myS.i)*	`4`	`int`	`0`
`-78` *(myS.firstC)*	`4`	`int`	`4`
`0` *(myS.secondC)*	`4`	`int`	`8`
`-78` *(mys.thirdC)*	`4`	`int`	`12`
`100` *(myS.j)*	`4`	`int`	`16`

Note that `myS.firstC` and `myS.thirdC` both use the instance ID `-78`. This allows the receiver to recognize that `firstC` and `thirdC` poin t at the same class instance (rather than at two different instances that happen to have the same contents).

Marshaling the negative instance IDs instead of the contents of an instance allows the receiver to accurately reconstruct the class graph that was sent by the sender. However, this begs the question of *when* the actual instances are to be marshaled as described at the beginning of this section. In Ice protocol messages, parameters and return values are marshaled as if they were members of a structure. For example, if an operation invocation has five input parameters, the client marshals the five parameters end-to-end as if they were members of a single structure. If any of the five parameters are class instances, or are of complex type (recursively) containing class instances, the sender marshals the parameters in multiple passes: the first pass marshals the parameters end-to-end, using the usual depth-first algorithm:

- If the sender encounters a class member during marshaling, it checks whether it has marshaled the same instance previously for the current request or reply:
  - If the instance has not been marshaled before, the sender assigns a new instance ID to the instance and marshals the negative ID.
  - Otherwise, if the instance was marshaled previously, the sender sends the same negative ID that it previously sent for that instance.

In effect, during marshaling, the sender builds a table that is indexed by the address of each instance; the lookup value for the instance is its ID.

Once the first pass ends, the sender has marshaled all the parameters, but has not yet marshaled any of the class instances that may be pointed at by various parameters or members. The instance table at this point contains all those instances for which negative IDs (pointers) were marshaled, so whatever is in the table at this point are the classes that the receiver still needs. The sender now marshals those instances in the table, but with positive IDs and followed by their contents, as described in our earlier example. The outstanding instances are marshaled as a sequence, that is, the sender marshals the number of instances as a size, followed by the actual instances.

In turn, the instances just sent may themselves contain class members; when those class members are marshaled, the sender assigns IDs to new instances or uses a negative ID for previously marshaled instances as usual. This means that, by the end of the second pass, the table may have grown, necessitating a third pass. That third pass again marshals the outstanding class instances as a size followed by the actual instances. The third pass contains all those instances that were not marshaled in the second pass. Of course, the third pass may trigger yet more passes until, finally, the sender has sent all outstanding instances, that is, marshaling is complete. At this point, the sender terminates the sequence of passes by marshaling an empty sequence (the value `0` encoded as a size).

To illustrate this with an example, consider the definitions shown in Classes with Operations once more:

Copyright 2016, ZeroC, Inc.

Slice

```
enum UnaryOp { UnaryPlus, UnaryMinus, Not };
enum BinaryOp { Plus, Minus, Multiply, Divide, And, Or };

class Node {
    idempotent long eval();
};

class UnaryOperator extends Node {
    UnaryOp operator;
    Node operand;
};

class BinaryOperator extends Node {
    BinaryOp op;
    Node operand1;
    Node operand2;
};

class Operand {
    long val;
};
```

These definitions allow us to construct expression trees. Suppose the client initializes a tree to the shape shown in the illustration below, representing the expression `(1 + 6 / 2) * (9 - 3)`. The values outside the nodes are the identities assigned by the client.

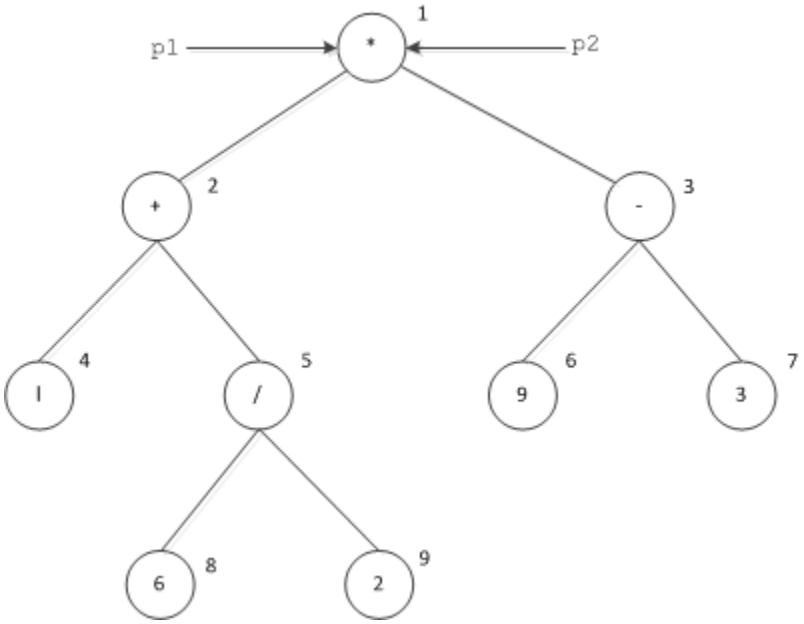

Expression tree for the expression `(1 + 6 / 2) * (9 - 3)`. Both p1 and p2 denote the root node.

The client passes the root of the tree to the following operation in the parameters `p1` and `p2`, as shown in the illustration above. (Even though it does not make sense to pass the same parameter value twice, we do it here for illustration purposes):

Copyright 2016, ZeroC, Inc.

Slice

```
interface Tree {
    void sendTree(Node p1, Node p2);
};
```

The client now marshals the two parameters `p1` and `p2` to the server, resulting in the value `-1` being sent twice in succession. (The client arbitrarily assigns an instance ID to each node. The value of the ID does not matter, as long as each node has a unique ID. For simplicity, the Ice implementation numbers instances with a counter that starts counting at `1` and increments by one for each unique instance.) This completes the marshaling of the parameters and results in a single instance with ID 1 in the instance table. The client now marshals a sequence containing a single element, node 1, as described in the example. In turn, node 1 results in nodes 2 and 3 being added to the table, so the next sequence of nodes contains two elements, nodes 2 and 3. The next sequence of nodes contains nodes 4, 5, 6, and 7, followed by another sequence containing nodes 8 and 9. At this point, no more class instances are outstanding, and the client marshals an empty sequence to indicate to the receiver that the final sequence has been marshaled.

Within each sequence, the order in which class instances are marshaled is irrelevant. For example, the third sequence could equally contain nodes 7, 6, 4, and 5, in that order. What is important here is that each sequence contains nodes that are an equal number of "hops" away from the initial node: the first sequence contains the initial node(s), the second sequence contains all nodes that can be reached by traversing a single link from the initial node(s), the third sequence contains all nodes that can be reached by traversing two links from the initial node(s), and so on.

Now consider the same example once more, but with different parameter values for `sendTree`: `p1` denotes the root of the tree, and `p2` deno tes the `-` operator of the right-hand sub-tree, as shown below:

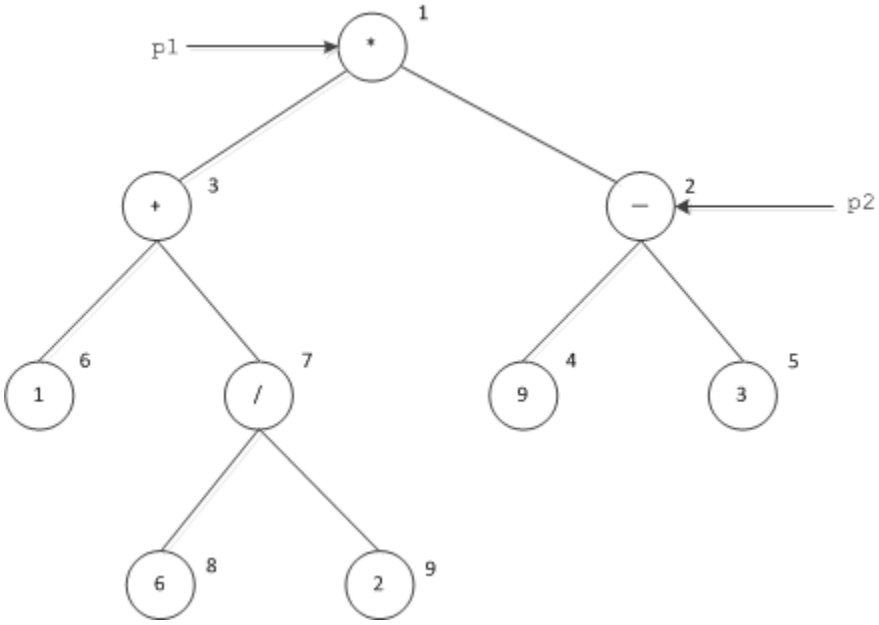

The expression tree of with p1 and p2 denoting different nodes.

The graph that is marshaled is exactly the same, but instances are marshaled in a different order and with different IDs:

- During the first pass, the client sends the IDs `-1` and `-2` for the parameter values.
- The second pass marshals a sequence containing nodes 1 and 2.
- The third pass marshals a sequence containing nodes 3, 4, and 5.
- The fourth pass marshals a sequence containing nodes 6 and 7.
- The fifth pass marshals a sequence containing nodes 8 and 9.
- The final pass marshals an empty sequence.

In this way, any graph of nodes can be transmitted (including graphs that contain cycles). The receiver reconstructs the graph by filling in a patch table during unmarshaling:

- Whenever the receiver unmarshals a negative ID, it adds that ID to a patch table; the lookup value is the memory address of the

Copyright 2016, ZeroC, Inc.

parameter or member that eventually will point at the corresponding instance.
- Whenever the receiver unmarshals an actual instance, it adds the instance to an unmarshaled table; the lookup value is the memory address of the instantiated class. The receiver then uses the address of the instance to patch any parameters or members with the actual memory address.

Note that the receiver may receive negative IDs that denote class instances that have been unmarshaled already (that is, point "backward" in the unmarshaling stream), as well as instances that are yet to be unmarshaled (that is, point "forward" in the unmarshaling stream). Both scenarios are possible, depending on the order in which instances are marshaled, as well as their in-degree.

To provide another example, consider the following definition:

Slice

```
class C {
    // ...
};

sequence<C> CSeq;
```

Suppose the client marshals a sequence of 100 `C` instances to the server, with each instance being distinct. (That is, the sequence contains 100 pointers to 100 different instances, not 100 pointers to the same single instance.) In that case, the sequence is marshaled as a size of 100, followed by 100 negative IDs, `-1` to `-100`. Following that, the client marshals a single sequence containing the 100 instances, each instance with its positive ID in the range `1` to `100`, and completes by marshaling an empty sequence.

On the other hand, if the client sends a sequence of 100 elements that all point to the same single class instance, the client marshals the sequence as a size of 100, followed by 100 negative IDs, all with the value `-1`. The client then marshals a sequence containing a single element, namely instance `1`, and completes by marshaling an empty sequence.

Encoding a Class Graph version 1.1

The most significant difference in the class encoding between version 1.0 and 1.1 is the location of class instances in the output stream. In version 1.0, instances are always marshaled at the end of the encapsulation, whereas in version 1.1 it is possible for instances to be marshaled *inline* at the point of first reference. Our previous discussion of class references describes the factors that determine whether a given class reference is marshaled as a reference or as an inline instance. The encoding rules can be summarized as follows:

- When using the compact format, always marshal an instance inline at the point of its first reference, and marshal all subsequent occurrences of the same instance as a reference.
- When using the sliced format, the encoding depends on the context in which a reference occurs: if the reference occurs while encoding a slice of an object or exception, encode it as an index into an indirection table that appears at the end of the slice, otherwise encode it as an inline instance or reference.

We can use the Slice definitions below to explore these situations:

Slice

```
class Node {
    int value;
    Node next;
};

struct S {
    Node obj;
};
```

Suppose we create an instance of structure `S` and assign it to the variable `s`, then construct the following class graph:

Copyright 2016, ZeroC, Inc.

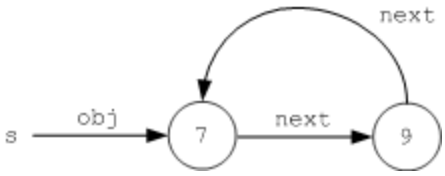

Class graph with circular reference.

We discuss the encoding for these values in separate sections below.

Class Graph in the Compact Format

As we marshal the structure's `obj` member, the compact format produces the following encoding:

Marshaled value	Size in bytes	Type	Byte offset
1 (*`obj` - inline instance marker - assigned ID 2*)	1	`size`	0
33 (*slice flags: string type ID, last slice*)	1	`byte`	1
`"::Node"` (*type ID - assigned index 1*)	7	`string`	2
7 (*`value`*)	4	`int`	9
1 (*`next` - inline instance marker - assigned ID 3*)	1	`size`	13
34 (*slice flags: type ID index, last slice*)	1	`byte`	14
1 (*type ID index*)	1	`size`	15
9 (*`value`*)	4	`int`	16
2 (*`next` - reference to ID 2*)	1	`size`	20

This example shows that inline instances are marshaled immediately upon discovery. As we encode the `obj` member, we encounter a new `Node` instance, assign it the ID 2, and begin to encode it inline using the reserved instance marker value `1`. Next, we encode the instance's `value` member and then examine its `next` member. Since we have not encountered this instance yet, we assign it the ID 3 and begin to encode this new instance inline. Its `next` member is a circular reference to the object with ID 2, so we simply encode a reference to this instance to complete the structure.

Class Graph in the Sliced Format

The need to support both slicing and slice preservation makes the sliced format more complex than the compact format. The encoding for our example values shows how class instances are handled differently depending on their context:

Marshaled value	Size in bytes	Type	Byte offset
1 (*`obj` - inline instance marker - assigned ID 2*)	1	`size`	0
57 (*slice flags: string type ID, size is present, indirection table is present, last slice*)	1	`byte`	1
`"::Node"` (*type ID - assigned index 1*)	7	`string`	2
9 (*byte count for slice*)	4	`int`	9
7 (*`value`*)	4	`int`	13
1 (*`next` - indirection table index*)	1	`size`	17
1 (*size of indirection table*)	1	`size`	18
1 (*indirection table entry - inline instance marker - assigned ID 3*)	1	`size`	19
58 (*slice flags: type ID index, size is present, indirection table is present, last slice*)	1	`byte`	20
1 (*type ID index*)	1	`size`	21

Copyright 2016, ZeroC, Inc.

9 (*byte count for slice*)	4	`int`	22
9 (*value*)	4	`int`	26
1 (*next - indirection table index - ID 2*)	1	`size`	30
1 (*size of indirection table*)	1	`size`	31
2 (*indirection table entry - reference to ID 2*)	1	`size`	32

For this example, let us assume that the structure value is an operation parameter, and therefore the structure member `obj` occurs *outside* t he context of a class or exception slice. As a result, we encode the instance with ID 2 as an inline instance. However, when we process the instance's `next` member, we are now *inside* the context of a class slice, so we must use an indirection table. The sender adds an entry to the indirection table for each unique class reference in the slice; instead of encoding a reference or inline instance, the sender encodes the index of the corresponding entry in the table. As shown above, the value encoded for the `next` member is `1`, representing the first entry in the table (`0` is still reserved for nil references).

The indirection table follows immediately after the last data member in the slice, but the table is *not* included in the byte count for the slice. A leading size denotes the number of entries in the table. The one and only entry in this table, the instance with ID 3, has not yet been encoded, therefore it is marshaled immediately as an inline instance. This instance also uses an indirection table, although in this case the table entry is simply a reference to instance ID 2, which has already been encoded.

Importance of the Indirection Table

The indirection table is necessary for implementing the slice preservation feature. Normally, when a receiver does not recognize the type ID in a slice, it has the option of ignoring that slice by skipping ahead in the stream by the number of bytes in the slice. However, when slice preservation is enabled, the receiver must keep a copy of the slice data in case the instance is later remarshaled. The need for the indirection table becomes apparent when you consider that an opaque blob of slice data may contain class references, and those class references can change during remarshaling. For example, without an indirection table, the sender might encode the instance ID 3 as the value of member `next`, but what happens if the receiver assigns that instance a different ID, such as `12`, when it remarshals the preserved slice? The receiver preserved the slice because it did not understand the type ID, which means it does not know the contents of the slice data and therefore it cannot "patch" any class references the slice might contain. The indirection table serves as an external "patch table" to solve this problem, essentially making the opaque slice data *relocatable* with respect to class references.

The byte count for a slice does not include the indirection table because the receiver must process the table regardless of whether it recognizes that slice's type ID. Consequently, to "skip" a slice, the receiver can skip (or preserve) the number of bytes specified by the slice's byte count, but still must decode the indirection table if the slice flags indicate that a table is present. If the receiver preserves the slice, it must also associate an indirection table with that slice; during remarshaling, the sender copies the opaque slice data into the stream, and then reconstructs the indirection table using (potentially) new instance IDs for the instances referenced in the table.

It is possible that the *only* reference to an instance is in an indirection table. To properly implement slice preservation, a receiver must therefore retain *every* instance that is referenced by a preserved indirection table. This is true even if the receiver does not recognize any of the type IDs in an instance; in effect, the receiver must construct a temporary "unknown object" placeholder for the instance, whose only purpose is to encapsulate the data comprising its slices in case the instance is later remarshaled.

Impact of Slicing on Class Graph Decoding

It is important to note that when a graph of class instances is sent, it always forms a connected graph. However, when the receiver rebuilds the graph, it may end up with a disconnected graph, due to slicing. Consider:

Copyright 2016, ZeroC, Inc.

Slice

```
class Base {
    // ...
};

class Derived extends Base {
    // ...
    Base b;
};

interface Example {
    void op(Base p);
};
```

Suppose the client has complete type knowledge, that is, understands both types `Base` and `Derived`, but the server only understands type `Base`, so the derived part of a `Derived` instance is sliced. The client can instantiate classes to be sent as parameter `p` as follows:

C++

```
DerivedPtr p = new Derived;
p->b = new Derived;
ExamplePrx e = ...;
e->op(p);
```

As far as the client is concerned, the graph looks like the one shown below:

Sender-side view of a graph containing derived instances.

However, the server does not understand the derived part of the instances and slices them. Yet, the server unmarshals all the class instances, leading to the situation where the class graph has become disconnected, as shown here:

Receiver-side view of the graph.

Of course, more complex situations are possible, such that the receiver ends up with multiple disconnected graphs, each containing many instances.

The slice preservation feature in version 1.1 of the encoding allows a receiver to remarshal the original graph intact, despite the fact that the receiver's in-memory object graph may appear to be disconnected.

See Also

- Classes with Operations
- Basic Data Encoding
- Data Encoding for Classes
- Simple Example of Class Encoding
- Protocol Messages
- Slicing Values and Exceptions

Copyright 2016, ZeroC, Inc.

Copyright 2016, ZeroC, Inc.

Data Encoding for Interfaces

Interfaces can be marshaled by value. For an interface marshaled by value (as opposed to a class instance derived from that interface), only the type ID of the most-derived interface is encoded. We use the following Slice definitions in our discussion:

Slice

```
interface Base { /* ... */ };

interface Derived extends Base { /* ... */ };

interface Example {
    void doSomething(Base b);
};
```

We assume that the client passes a class instance to `doSomething` that does not have a Slice definition (but implements `Derived`).

Interface Encoding version 1.0

The on-the-wire representation of the interface parameter using encoding version 1.0 is as follows:

Marshaled Value	Size in Bytes	Type	Byte offset
`1` *(identity)*	4	`int`	0
`0` *(marker for class type ID)*	1	`bool`	4
`"::Derived"` *(class type ID)*	10	`string`	5
`4` *(byte count for slice)*	4	`int`	15
`0` *(marker for class type ID)*	1	`bool`	19
`"::Ice::Object"` *(class type ID)*	14	`string`	20
`5` *(byte count for slice)*	4	`int`	34
`0` *(number of dictionary entries)*	1	`size`	38

Interface Encoding version 1.1

With encoding version 1.1, the interface parameter is encoded as if it was a class with no data members. Here is the on-the-wire representation using the sliced format:

Marshaled value	Size in bytes	Type	Byte offset
`1` *(instance marker)*	1	`size`	0
`17` *(slice flags: string type ID, size is present)*	1	`byte`	1
`"::Derived"` *(type ID - assigned index 1)*	10	`string`	2
`4` *(byte count for slice)*	4	`int`	12

The compact format omits the byte count for the slice:

Marshaled value	Size in bytes	Type	Byte offset
`1` *(instance marker)*	1	`size`	0
`1` *(slice flags: string type ID)*	1	`byte`	1
`"::Derived"` *(type ID - assigned index 1)*	10	`string`	2

Copyright 2016, ZeroC, Inc.

See Also

- Passing Interfaces by Value
- Type IDs
- Data Encoding for Classes

Copyright 2016, ZeroC, Inc.

Data Encoding for Proxies

On this page:

- Encoding for General Proxy Parameters
- Encoding for Endpoint Parameters
- Encoding for TCP Endpoint Parameters
- Encoding for UDP Endpoint Parameters
- Encoding for SSL Endpoint Parameters
- Encoding for WS Endpoint Parameters
- Encoding for WSS Endpoint Parameters

Encoding for General Proxy Parameters

The encoding format of proxies changed in version 1.1.

Proxy Parameter Encoding version 1.0

The first component of an encoded proxy is a value of type `Ice::Identity`. If the proxy is a nil value, the `category` and `name` members are empty strings, and no additional data is encoded. The encoding for a non-null proxy consists of general parameters followed by endpoint parameters.

The general proxy parameters are encoded as if they were members of the following structure:

Slice

```
struct ProxyData {
    Ice::Identity id;
    Ice::StringSeq facet;
    byte mode;
    bool secure;
};
```

The general proxy parameters are described in the table below:

Parameter	Description
`id`	The object identity
`facet`	The facet name (zero- or one-element sequence)
`mode`	The proxy mode (`0`=twoway, `1`=oneway, `2`=batch oneway, `3`=datagram, `4`=batch datagram)
`secure`	`true` if secure endpoints are required, otherwise `false`

The `facet` field has either zero elements or one element. An empty sequence denotes the default facet, and a one-element sequence provides the facet name in its first member. If a receiver receives a proxy with a `facet` field with more than one element, it must throw a `ProxyUnmarshalException`.

Proxy Parameter Encoding version 1.1

Version 1.1 of the encoding appends two attributes to the existing proxy parameters in version 1.0: protocol and encoding versions. The general proxy parameters are encoded as if they were members of the following structure:

Copyright 2016, ZeroC, Inc.

Slice

```
struct ProtocolVersion {
    byte major;
    byte minor;
};

struct EncodingVersion {
    byte major;
    byte minor;
};
struct ProxyData {
    Ice::Identity id;
    Ice::StringSeq facet;
    byte mode;
    bool secure;
    ProtocolVersion protocol;
    EncodingVersion encoding;
};
```

The additional parameters are described in the table below:

Parameter	Description
`protocol`	The maximum protocol version supported by the server. Currently this value is always 1.0.
`encoding`	The maximum encoding version supported by the server.

The encoding for UDP endpoints also changed in version 1.1.

Encoding for Endpoint Parameters

A proxy optionally contains an endpoint list or an adapter identifier, but not both:

- If a proxy contains endpoints, they are encoded immediately following the general parameters. A size specifying the number of endpoints is encoded first, followed by the endpoints. Each endpoint is encoded as a `short` specifying the endpoint type (`1`=TCP, `2`=SSL, 3=UDP, 4=WS, 5=WSS), followed by an encapsulation of type-specific parameters. The type-specific parameters for TCP, UDP, SSL, WS and WSS are presented in the sections that follow.
- If a proxy does not have endpoints, a single byte with value `0` immediately follows the general parameters and a string representing the object adapter identifier is encoded immediately following the zero byte.
- For a proxy to a well-known object, which has neither endpoints nor an object adapter identifier, a single byte with value `0` immediat ely follows the general parameters and an empty string is encoded immediately following the zero byte.

Type-specific endpoint parameters are encapsulated because a receiver may not be capable of decoding them. For example, a receiver can only decode SSL endpoint parameters if it is configured with the IceSSL plug-in. However, the receiver must be able to re-encode the proxy with all of its original endpoints, in the order they were received, even if the receiver does not understand the type-specific parameters for an endpoint. Encapsulation of the parameters into an opaque endpoint allows the receiver to do this.

Encoding for TCP Endpoint Parameters

A TCP endpoint is encoded as an encapsulation containing the following structure:

Copyright 2016, ZeroC, Inc.

Slice

```
struct TCPEndpointData {
    string host;
    int port;
    int timeout;
    bool compress;
};
```

The endpoint parameters are described in the following table.

Parameter	Description
`host`	The server host (a host name or IP address)
`port`	The server port (1-65535)
`timeout`	The timeout in milliseconds for socket operations
`compress`	`true` if compression should be used (if possible), otherwise `false`

Encoding for UDP Endpoint Parameters

The encoding format of UDP endpoints changed in version 1.1.

UDP Endpoint Encoding version 1.0

A UDP endpoint is encoded as an encapsulation containing the following structure:

Slice

```
struct UDPEndpointData {
    string host;
    int port;
    byte protocolMajor;
    byte protocolMinor;
    byte encodingMajor;
    byte encodingMinor;
    bool compress;
};
```

The endpoint parameters are described in the following table.

Parameter	Description
`host`	The server host (a host name or IP address)
`port`	The server port (1-65535)
`protocolMajor`	The major protocol version supported by the endpoint
`protocolMinor`	The highest minor protocol version supported by the endpoint

Copyright 2016, ZeroC, Inc.

`encodingMajor`	The major encoding version supported by the endpoint
`encodingMinor`	The highest minor encoding version supported by the endpoint
`compress`	`true` if compression should be used (if possible), otherwise `false`

UDP Endpoint Encoding version 1.1

Version 1.1 of the encoding omits the protocol and encoding versions because these attributes are handled as general proxy parameters instead:

Slice

```
struct UDPEndpointData {
    string host;
    int port;
    bool compress;
};
```

The endpoint parameters are described in the following table.

Parameter	Description
`host`	The server host (a host name or IP address)
`port`	The server port (1-65535)
`compress`	`true` if compression should be used (if possible), otherwise `false`

Encoding for SSL Endpoint Parameters

An SSL endpoint is encoded as an encapsulation containing the following structure:

Slice

```
struct SSLEndpointData {
    string host;
    int port;
    int timeout;
    bool compress;
};
```

The endpoint parameters are described in the following table.

Parameter	Description
`host`	The server host (a host name or IP address)
`port`	The server port (1-65535)
`timeout`	The timeout in milliseconds for socket operations

Copyright 2016, ZeroC, Inc.

`compress`	`true` if compression should be used (if possible), otherwise `false`

Encoding for WS Endpoint Parameters

A WebSocket endpoint is encoded as an encapsulation containing the following structure:

Slice

```
struct WSEndpointData {
    string host;
    int port;
    int timeout;
    bool compress;
    string resource;
};
```

The endpoint parameters are described in the following table.

Parameter	Description
`host`	The server host (a host name or IP address)
`port`	The server port (1-65535)
`timeout`	The timeout in milliseconds for socket operations
`compress`	`true` if compression should be used (if possible), otherwise `false`
`resource`	A URI representing the web server resource associated with this endpoint

Encoding for WSS Endpoint Parameters

A WebSocket endpoint is encoded as an encapsulation containing the following structure:

Slice

```
struct WSSEndpointData {
    string host;
    int port;
    int timeout;
    bool compress;
    string resource;
};
```

The endpoint parameters are described in the following table.

Parameter	Description
`host`	The server host (a host name or IP address)
`port`	The server port (1-65535)
`timeout`	The timeout in milliseconds for socket operations

Copyright 2016, ZeroC, Inc.

`compress`	`true` if compression should be used (if possible), otherwise `false`
`resource`	A URI representing the web server resource associated with this endpoint

See Also

- Object Identity
- Versioning
- Basic Data Encoding
- IceSSL
- Using Connections
- Protocol Compression

Copyright 2016, ZeroC, Inc.

Data Encoding for Optional Values

Slice definitions can use optional values as parameters and as data members. To make optional values truly useful, the receiver needs the ability to gracefully ignore an optional value that it does not recognize. Although the Ice encoding omits type information and therefore is not self-describing, the encoding for optional values must include enough information for a receiver to determine how many bytes an optional value occupies so that it can skip to the next value in the stream. Despite this requirement, the encoding rules minimize the overhead associated with optional values, as you will see below.

On this page:

- Overview of the Optional Value Encoding
- Encoding for Optional Types and Tags
- Examples of Optional Value Encoding
  - Optional Parameters Example
  - Optional Data Members Example

Overview of the Optional Value Encoding

The encoding for optional parameters and data members follows these general rules:

- Each optional value has a corresponding integer tag that uniquely identifies it.
- Optional values always appear *after* any required values.
- Optional values must be sorted by their tags. (Unlike required values, order of declaration does not affect optional values.)
- An optional value is encoded only if the sender has supplied a value.

Optional data members of a class or exception appear in the slice after any required data members and before the indirection table, if present. The slice flags must indicate the presence of optional data members, which are included in the byte count for the slice. If a slice contains optional data members, the byte value 255 must be written after the last optional data member; this marker denotes the end of the optional data members and is also included in the byte count for the slice.

The encoding does not use an explicit end marker for optional parameters; the end of the encapsulation also marks the end of any optional parameters.

An optional value is encoded as the tuple *<type, tag, value>*, where *type* is the *optional type* that tells the receiver how to determine the number of bytes occupied by the value. The value itself is marshaled using the standard Ice encoding rules for its Slice type.

Optional values require encoding version 1.1.

Encoding for Optional Types and Tags

The first byte of an encoded optional value includes the optional type, and may also include the tag. The optional type occupies the first three bits of this byte, as described in the table below:

Name	Value	Description	Used for Slice type
F1	0	The value is encoded in one byte.	bool, byte
F2	1	The value is encoded in two bytes.	short
F4	2	The value is encoded in four bytes.	int, float
F8	3	The value is encoded in eight bytes.	double, long
Size	4	The value is encoded as a size.	enum
VSize	5	A leading size value indicates the number of bytes occupied by the value.	string, fixed-size structure, container of fixed-size elements
FSize	6	A leading 32-bit integer indicates the number of bytes occupied by the value.	variable-size structure, container of variable-size elements
Class	7	A class reference or inline instance.	class

The next five bits of the leading byte contain the tag, but only if the tag value is less than 30. Otherwise, the next five bits contain the value 30 as a marker to indicate that the tag value is encoded as a size starting with the next byte. As you can see, using tag values in the range 0 to 29 produces the most compact encoding.

Variable-size types whose encoded size cannot be determined in advance use the FSize optional type, where "FSize" denotes a leading fixed-length (32-bit) size. For these types, the sender reserves four bytes to hold the size, encodes the value as usual, then replaces the four

Copyright 2016, ZeroC, Inc.

bytes with the actual encoded size. This strategy avoids the need to shift the encoded data in the buffer, at the expense of potentially consuming more bytes than necessary to encode the size.

Fixed-size types use the VSize optional type, where "VSize" denotes a leading variable-length size. The sender can determine the encoded size of these types in advance, and therefore encodes it as a size followed by the value as usual.

Strings also use the VSize optional type but do not require an additional size because the string encoding already includes a leading size. The same is true for sequences of `bool` and `byte`.

The optional type Class represents a class reference or inline instance. Note that a receiver must decode an inline instance even if it does not recognize the tag value because the instance may be referenced by other parameters.

The following table describes the encoding of Slice types:

Slice type	Optional type	Data encoding	Notes
`bool`, `byte`	F1	value	
`short`	F2	value	
`int`, `float`	F4	value	
`long`, `double`	F8	value	
Proxy	FSize	int + data	32-bit integer holds the size of the encoded proxy
`class`	Class	reference or instance	
`enum`	Size	size	Enumerator encoded as a size
`string`	VSize	value	The encoded data for the string already contains a leading size
`sequence<bool>`, `sequence<byte>`	VSize	value	The encoded data for the sequence already contains a leading size
`sequence<`*fixed-size type*`>`, `dictionary<`*fixed-size key*,*fixed-size value*`>`	VSize	size + value	Size can be computed before encoding the container
`sequence<`*variable size type*`>`, `dictionary<`*variable-size key*,*variable-size value*`>`	FSize	int + value	32-bit integer holds the size of the container

Examples of Optional Value Encoding

The examples presented below demonstrate the encoding for optional values with typical use cases.

Optional Parameters Example

The following Slice operation makes use of optional input and output parameters:

Slice

```
bool op1(byte b,
         optional(2) string name,
         short sh,
         optional(1) long count,
         out double d,
         out optional(300) Object* p);
```

Suppose the parameters have the values shown in the table below:

Member	Type	Value	Marshaled size (in bytes)
`b`	`byte`	`77`	`1`

Copyright 2016, ZeroC, Inc.

`name`	`string`	`"joe"`	4
`sh`	`short`	`99`	2
`count`	`long`	`88`	8
`d`	`double`	`3.14`	8
`p`	`proxy`	nil	2
return value	`bool`	`true`	1

The parameters of the outgoing request are encoded in an encapsulation with all required parameters first, in order of declaration, followed by the optional parameters sorted by tag:

Marshaled value	Size in bytes	Type	Byte offset
`77` (*b*)	1	`byte`	0
`99` (*sh*)	2	`short`	1
`11` (*optional type F8 + tag 1*)	1	`byte`	3
`88` (*count*)	8	`long`	4
`21` (*optional type VSize + tag 2*)	1	`byte`	12
`"joe"` (*name*)	4	`long`	13

Using tag values less than 30 allows the encoding to combine the optional type and the tag into the same byte using the expression `(Tag << 3) + Type`.

Now consider the contents of the reply message:

Marshaled value	Size in bytes	Type	Byte offset
`3.14` (*d*)	8	`double`	0
`true` (*return value*)	1	`bool`	8
`246` (*optional type FSize + marker 30*)	1	`byte`	9
`300` (*tag*)	5	`size`	10
`2` (*32-bit FSize*)	4	`int`	15
nil (*p*)	2	`Object*`	19

The body of the reply message contains the out parameter `d`, the return value, and the optional parameter `p`. The tag value 300 is too large to combine with the optional type, therefore it appears immediately following the optional type encoded as a size. A proxy value uses the FSize optional type, meaning a 32-bit integer precedes the encoded value to specify its size.

Although the return value is required in this example, an optional return value is treated as if it were an optional out parameter.

The server here supplies a nil value for the optional proxy parameter. The client must not interpret this to mean that the optional parameter is unset; rather, the parameter is set, it just happens to be set to a nil value. If the server had supplied no value for the parameter, it would not appear in the encoding at all.

As an example, if the client does not recognize the optional proxy parameter, it can skip the parameter as follows:

1. Extract the byte containing the optional type
2. Examine the first three bits to determine the type
3. Examine the five remaining bits to determine the tag's status
4. A value of 30 means the tag is encoded separately
5. Extract the tag as a size
6. The optional type FSize means a 32-bit integer holds the value's size; extract the integer
7. Skip ahead the specified number of bytes
8. If we have reached the end of the encapsulation, there are no more optional parameters remaining

Optional Data Members Example

Copyright 2016, ZeroC, Inc.

The only difference between the encoding for optional parameters and the encoding for optional data members is the latter uses a marker byte to signify the end of optional parameters in a slice. Consider the following definitions:

Slice

```
struct Color
{
    short red;
    short green;
    short blue;
};

class Shape
{
    optional(1) string label;
};

class Rectangle extends Shape
{
    int width;
    int height;
    optional(10) Color fill;
    optional(9) Color border;
    optional(11) float scale;
};
```

Suppose we are encoding an instance of `Rectangle` with the member values shown in the table below:

Member	Type	Value	Marshaled size (in bytes)
`label`	`string`	`"r1"`	3
`width`	`int`	`41`	4
`height`	`int`	`16`	4
`fill`	`Color`	`0,0,0`	6
`border`	`Color`	`255,255,255`	6
`scale`	`float`	`2.0`	4

Using the sliced format, the instance data looks as follows:

Marshaled value	Size in bytes	Type	Byte offset
`1` (*instance marker*)	1	`byte`	0
`21` (*slice flags: string type ID, size is present, optional members are present*)	1	`byte`	1
`"::Rectangle"` (*type ID*)	12	`string`	2
`34` (*byte count for slice*)	4	`int`	14
`41` (*`width`*)	4	`int`	18
`16` (*`height`*)	4	`int`	22
`77` (*optional type VSize + tag 9*)	1	`byte`	26
`6` (*VSize*)	1	`size`	27

Copyright 2016, ZeroC, Inc.

`{0,0,0}` (*`border`*)	`6`	`Color`	`28`
`85` (*optional type VSize + tag `10`*)	`1`	`byte`	`34`
`6` (*VSize*)	`1`	`size`	`35`
`{255,255,255}` (*`fill`*)	`6`	`Color`	`36`
`90` (*optional type F4 + tag `11`*)	`1`	`byte`	`42`
`2.0` (*`scale`*)	`4`	`float`	`43`
`255` (*end marker*)	`1`	`byte`	`47`
`53` (*slice flags: string type ID, size is present, optional members are present, last slice*)	`1`	`byte`	`48`
`"::Shape"` (*type ID*)	`8`	`string`	`49`
`9` (*byte count for slice*)	`4`	`int`	`57`
`13` (*optional type VSize + tag `1`*)	`1`	`byte`	`61`
`"r1"` (*`label`*)	`3`	`string`	`62`
`255` (*end marker*)	`1`	`byte`	`65`

Notice the use of an end marker (byte value `255`) denoting the end of the optional data members in each slice.

See Also

- Optional Values
- Basic Data Encoding
- Data Encoding for Classes

Copyright 2016, ZeroC, Inc.

Protocol Messages

The Ice protocol uses five messages:

- Request (from client to server)
- Batch request (from client to server)
- Reply (from server to client)
- Validate connection (from server to client)
- Close connection (client to server or server to client)

Of these messages, validate and close connection only apply to connection-oriented transports.

As with the data encoding, protocol messages have no alignment restrictions. Each message consists of a message header and (except for validate and close connection) a message body that immediately follows the header.

On this page:

- Message Header
- Request Message Body
- Batch Request Message Body
- Reply Message Body
- Validate Connection Message
- Close Connection Message
- Protocol State Machine
- Disorderly Connection Closure

Message Header

Each protocol message has a 14-byte header that is encoded as if it were the following structure:

Slice

```
struct HeaderData {
    int  magic;
    byte protocolMajor;
    byte protocolMinor;
    byte encodingMajor;
    byte encodingMinor;
    byte messageType;
    byte compressionStatus;
    int  messageSize;
};
```

The message header members are described in the following table.

Member	Description
`magic`	A four-byte magic number consisting of the ASCII-encoded values of 'I', 'c', 'e', 'P' (0x49, 0x63, 0x65, 0x50)
`protocolMajor`	The protocol major version number
`protocolMinor`	The protocol minor version number
`encodingMajor`	The encoding major version number
`encodingMinor`	The encoding minor version number
`messageType`	The message type
`compressionStatus`	The compression status of the message

Copyright 2016, ZeroC, Inc.

`messageSize`	The size of the message in bytes, including the header

Currently, both the protocol and the encoding are at version 1.0. The valid message types are shown in the following table.

Message Type	Encoding
Request	`0`
Batch request	`1`
Reply	`2`
Validate connection	`3`
Close connection	`4`

The encoding for the message bodies of each of these message types is described in the sections that follow.

Request Message Body

A request message contains the data necessary to perform an invocation on an object, including the identity of the object, the operation name, and input parameters. A request message is encoded as if it were the following structure:

Slice

```
struct RequestData {
    int requestId;
    Ice::Identity id;
    Ice::StringSeq facet;
    string operation;
    byte mode;
    Ice::Context context;
    Encapsulation params;
};
```

The request members are described in the following table.

Member	Description
`requestId`	The request identifier
`id`	The object identity
`facet`	The facet name (zero- or one-element sequence)
`operation`	The operation name
`mode`	A byte representation of `Ice::OperationMode` (`0`=normal, `2`=idempotent)
`context`	The invocation context
`params`	The encapsulated input parameters, in order of declaration

The request identifier zero (`0`) is reserved for use in oneway requests and indicates that the server must not send a reply to the client. A non-zero request identifier must uniquely identify the request on a connection, and must not be reused while a reply for the identifier is outstanding.

The `facet` field has either zero elements or one element. An empty sequence denotes the default facet, and a one-element sequence provides the facet name in its first member. If a receiver receives a request with a `facet` field with more than one element, it must throw a `M`

Copyright 2016, ZeroC, Inc.

`arshalException`.

Batch Request Message Body

A batch request message contains one or more oneway requests, bundled together for the sake of efficiency. A batch request message is encoded as integer (not a size) that specifies the number of requests in the batch, followed by the corresponding number of requests, encoded as if each request were the following structure:

Slice

```
struct BatchRequestData {
    Ice::Identity id;
    Ice::StringSeq facet;
    string operation;
    byte mode;
    Ice::Context context;
    Encapsulation params;
};
```

The batch request members are described in the following table.

Member	Description
`id`	The object identity
`facet`	The facet name (zero- or one-element sequence)
`operation`	The operation name
`mode`	A byte representation of `Ice::OperationMode`
`context`	The invocation context
`params`	The encapsulated input parameters, in order of declaration

Note that no request ID is necessary for batch requests because only oneway invocations can be batched.

The `facet` field has either zero elements or one element. An empty sequence denotes the default facet, and a one-element sequence provides the facet name in its first member. If a receiver receives a batch request with a `facet` field with more than one element, it must throw a `MarshalException`.

Reply Message Body

A reply message body contains the results of a twoway invocation, including any return value, out-parameters, or exception. A reply message body is encoded as if it were the following structure:

```
struct ReplyData {
    int requestId;
    byte replyStatus;
    Encapsulation body; // messageSize - 19 bytes
};
```

The first four bytes of a reply message body contain a request ID. The request ID matches an outgoing request and allows the requester to associate the reply with the original request.

Copyright 2016, ZeroC, Inc.

The byte following the request ID indicates the status of the request; the remainder of the reply message body following the status byte is an encapsulation whose contents depend on the status value. The possible reply status values are shown in the table below (most of these values correspond to common exceptions).

<table>
<tr><th>Reply status</th><th>Success</th><th>Description</th></tr>
<tr><td>Success</td><td>0</td><td>A successful reply message is encoded as an encapsulation containing out-parameters (in the order of declaration), followed by the return value for the invocation, encoded according to their types as specified in Data Encoding. If an operation declares a <code>void</code> return type and no out-parameters, an empty encapsulation is encoded.</td></tr>
<tr><td>User exception</td><td>1</td><td>A user exception reply message contains an encapsulation containing the encoded user exception.</td></tr>
<tr><td>Object does not exist</td><td>2</td><td>If the target object does not exist, the reply message is encoded as if it were the following structure inside an encapsulation:
Slice
<pre>struct ReplyData {
    Ice::Identity id;
    Ice::StringSeq facet;
    string operation;
};</pre>
The invalid object reply members are described below:
<table>
<tr><th>Member</th><th>Description</th></tr>
<tr><td><code>id</code></td><td>object identity</td></tr>
<tr><td><code>facet</code></td><td>facet</td></tr>
<tr><td><code>operation</code></td><td>The operation name</td></tr>
</table>
The <code>facet</code> field has either zero elements or one element. An empty sequence denotes the default facet, and a one-element sequence provides the facet name in its first member. If a receiver receives a reply with a <code>facet</code> field with more than one element, it must throw a <code>MarshalException</code>.</td></tr>
<tr><td>Facet does not exist</td><td>3</td><td>If the target object does not support the facet encoded in the request message, the reply message is encoded as for reply status 2.</td></tr>
<tr><td>Operation does not exist</td><td>4</td><td>If the target object does not support the operation encoded in the request message, the reply message is encoded as for reply status 2.</td></tr>
<tr><td>Unknown Ice local exception</td><td>5</td><td>The reply message for an unknown Ice local exception is encoded as an encapsulation containing a single string that describes the exception.</td></tr>
<tr><td>Unknown Ice user exception</td><td>6</td><td>The reply message for an unknown Ice user exception is encoded as an encapsulation containing a single string that describes the exception.</td></tr>
<tr><td>Unknown exception</td><td>7</td><td>The reply message for an unknown exception is encoded as an encapsulation containing a single string that describes the exception.</td></tr>
</table>

Validate Connection Message

A server sends a validate connection message when it receives a new connection.

Validate connection messages are only used for connection-oriented transports.

The message indicates that the server is ready to receive requests; the client must not send any messages on the connection until it has

Copyright 2016, ZeroC, Inc.

received the validate connection message from the server. No reply to the message is expected by the server.

The purpose of the validate connection message is two-fold:

- It confirms to the client that the server is indeed an Ice server and not another type of server reached by accident.
- It prevents the client from writing a request message to its local transport buffers until after the server has acknowledged that it can actually process the request. This avoids a race condition caused by the server's TCP/IP stack accepting connections in its backlog while the server is in the process of shutting down: if the client were to send a request in this situation, the request would be lost but the client could not safely re-issue the request because that might violate at-most-once semantics. The validate connection message guarantees that a server is not in the middle of shutting down when the server's TCP/IP stack accepts an incoming connection and so avoids the race condition.

As of Ice 3.6, validate connection messages may also be sent at any time by either side as a heartbeat. See Active Connection Management for more information.

The message header comprises the entire validate connection message. The compression status of a validate connection message is always `0`.

Close Connection Message

A close connection message is sent when a peer is about to gracefully shutdown a connection.

Close connection messages are only used for connection-oriented transports.

The message header comprises the entire close connection message. The compression status of a close connection message is always `0`.

Either client or server can initiate connection closure. On the client side, connection closure is triggered by Active Connection Management (ACM), which automatically reclaims connections that have been idle for some time.

This means that connection closure can be initiated at will by either end of a connection; most importantly, no state is associated with a connection as far as the object model or application semantics are concerned.

The client side can close a connection whenever no reply for a request is outstanding on the connection. The sequence of events is:

1. The client sends a close connection message.
2. The client closes the writing end of the connection.
3. The server responds to the client's close connection message by closing the connection.

The server side can close a connection whenever no operation invocation is in progress that was invoked via that connection. This guarantees that the server will not violate at-most-once semantics: an operation, once invoked in a servant, is allowed to complete and its results are returned to the client. Note that the server can close a connection even after it has received a request from the client, provided that the request has not yet been passed to a servant. In other words, if the server decides that it wants to close a connection, the sequence of events is:

1. The server discards all incoming requests on the connection.
2. The server waits until all still executing requests have completed and their results have been returned to the client.
3. The server sends a close connection message to the client.
4. The server closes its writing end of the connection.
5. The client responds to the server's close connection message by closing both its reading and writing ends of the connection.
6. If the client has outstanding requests at the time it receives the close connection message, it re-issues these requests on a new connection. Doing so is guaranteed not to violate at-most-once semantics because the server guarantees not to close a connection while requests are still in progress on the server side.

Protocol State Machine

From a client's perspective, the Ice protocol behaves according to the state machine shown below:

Copyright 2016, ZeroC, Inc.

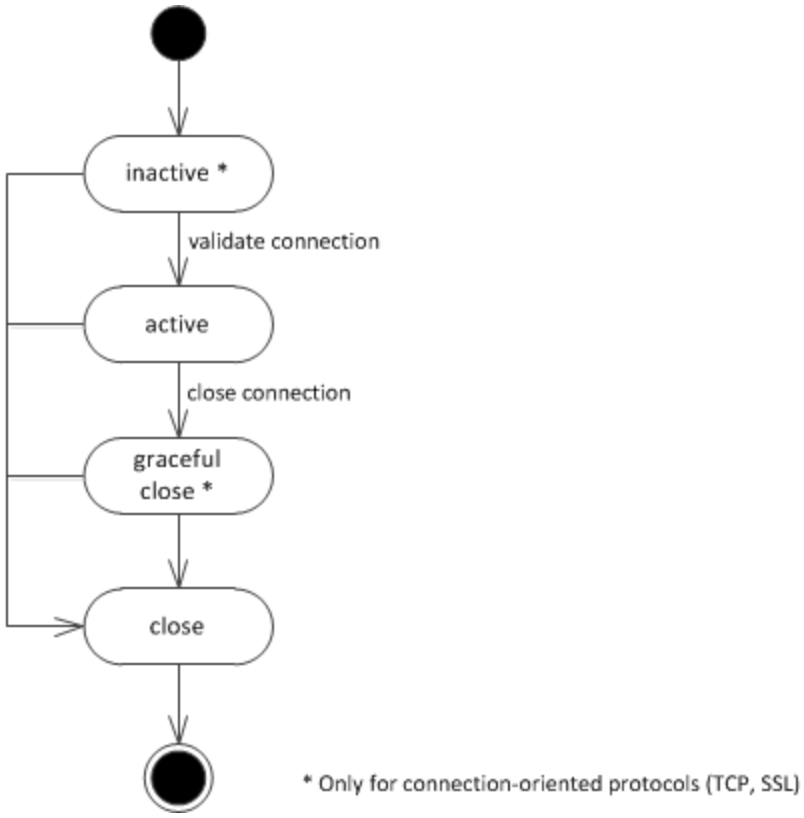

Protocol state machine.

To summarize, a new connection is inactive until a validate connection message has been received by the client, at which point the active state is entered. The connection remains in the active state until it is shut down, which can occur when there are no more proxies using the connection, or after the connection has been idle for a while. At this point, the connection is gracefully closed, meaning that a close connection message is sent, and the connection is closed.

Disorderly Connection Closure

Any violation of the protocol or encoding rules results in a disorderly connection closure: the side of the connection that detects a violation unceremoniously closes it (without sending a close connection message or similar). There are many potential error conditions that can lead to disorderly connection closure; for example, the receiver might detect that a message has a bad magic number or incompatible version, receive a reply with an ID that does not match that of an outstanding request, receive a validate connection message when it should not, or find illegal data in a request (such as a negative size, or a size that disagrees with the actual data that was unmarshaled).

See Also

- Data Encoding
- Protocol Compression
- Object Identity
- Versioning
- Request Contexts
- Oneway Invocations
- Batched Invocations
- Protocol and Encoding Versions
- Active Connection Management
- Automatic Retries
- Connection Closure

Copyright 2016, ZeroC, Inc.

Protocol Compression

Compression is an optional feature of the Ice protocol; whether it is used for a particular message is determined by several factors:

1. Compression may not be supported on all platforms or in all language mappings.
2. Compression can be used in a request or batch request only if the endpoint advertises the ability to accept compressed messages.
3. For efficiency reasons, the Ice protocol engine does not compress messages smaller than 100 bytes.

A compliant implementation of the protocol is free to compress messages that are smaller than 100 bytes — the choice is up to the protocol implementation.

If compression is used, the entire message excluding the header is compressed using the bzip2 algorithm [1]. The `messageSize` member of the message header therefore reflects the size of the compressed message, including the uncompressed header, plus an additional four bytes.

The `compressionStatus` field of the message header indicates whether a message is compressed and whether the sender can accept a compressed reply, as shown in the table below.

Reply status	Encoding	Applies to	Description
Message is uncompressed, sender cannot accept a compressed reply.	`0`	Request, Batch Request, Reply, Validate Connection, Close Connection	A client that does not support compression always uses this value. A client that supports compression sets the value to `0` if the endpoint via which the request is dispatched indicates that it does not support compression. A server uses this value for uncompressed replies.
Message is uncompressed, sender can accept a compressed reply.	`1`	Request, Batch Request	A client uses this value if the endpoint via which the request is dispatched indicates that it supports compression, but the client has decided not to use compression for this particular request (presumably because the request is too small, so compression does not provide any saving).
Message is compressed and sender can accept a compressed reply.	`2`	Request, Batch Request, Reply	A client that supports compression sets this value only if the endpoint via which the request is dispatched indicates that it supports compression. A server uses this value for compressed replies.

The message body of a compressed request, batch request, or reply message is encoded by first writing the size of the uncompressed message (including its header) as a four-byte integer, followed by the compressed message body (excluding the header). It follows that the size of a compressed message is 14 bytes for the header, plus four bytes to record the size of the uncompressed message, plus the number of bytes occupied by the compressed message body. Writing the uncompressed message size prior to the body enables the receiver to allocate a buffer that is large enough to accomodate the uncompressed message body.

Note that compression is likely to improve performance only over lower-speed links, for which bandwidth is the overall limiting factor. Over high-speed LAN links, the CPU time spent on compressing and uncompressing messages is longer than the time it takes to just send the uncompressed data.

See Also

- Data Encoding for Proxies
- Protocol Messages

References

1. Red Hat, Inc. 2003. *The bzip2 and libbzip2 Home Page*. Raleigh, NC: Red Hat, Inc.

Copyright 2016, ZeroC, Inc.

Protocol and Encoding Versions

On this page:

- Version Flexibility
- Version Ground Rules

Version Flexibility

As we saw in the preceding sections, both the Ice protocol and encoding have separate major and minor version numbers. Separate versioning of protocol and encoding has the advantage that neither depends on the other: any version of the Ice protocol can be used with any version of the encoding, so they can evolve independently. For example, you could send a request with Ice protocol version 1.1, and encode the parameters with encoding version 2.3.

The Ice versioning mechanism provides the maximum possible amount of interoperability between clients and servers that use different versions of the Ice run time. In particular, older deployed clients can communicate with newer deployed servers and vice versa, provided that the message contents use types that are understandable to both sides. Intermediary servers, such as a Glacier2 router, do not and should not check the contents of the requests and responses they forward; this way, an older intermediary server can forward messages encoded with a newer encoding that it does not understand.

Due to a bug, a Glacier2 router prior to version 3.5 can only forward messages encoded using the 1.0 encoding.

Ice 3.5 introduced a new encoding, encoding version 1.1, in support of several new features such as the ability to define optional parameters in operations. In order to send a request with an optional parameter, you need to use encoding 1.1 or greater.

Ice 3.5 (or newer) is capable of marshaling and unmarshaling messages with both the 1.0 and 1.1 encodings (it uses the 1.1 encoding by default). This allows interoperability to the maximum extent possible:

Client Ice Version	Server Ice Version	Encoding Version	Operation without Optional Parameter	Operation with Optional Parameter
3.4 or earlier	3.4 or earlier	1.0	Yes	No
3.4 or earlier	3.5 or newer	1.0	Yes	No
3.5 or newer	3.4 or earlier	1.0	Yes	No
3.5 or newer	3.5 or newer	1.1	Yes	Yes

When an Ice 3.5 (or newer) client sends a request to an Ice 3.4 server, it must naturally use the 1.0 encoding. This is achieved by setting correctly (and often automatically) the encoding version on the proxy used. If this proxy is received by the client through a response using the 1.0 encoding, the decoded proxy will automatically have the desired 1.0 encoding (all proxies in a 1.0-encoded message implicitly have 1.0 as their encoding version and 1.0 as their protocol version). If this proxy is created with `stringToProxy` or `propertyToProxy`, the client will need to set this proxy's encoding version to 1.0 through one of the following methods:

- with `-e 1.0` in the source stringified proxy
- by creating a new proxy with the desired encoding version, by calling `ice_encodingVersion("1.0")` on the proxy
- by setting `Ice.Default.EncodingVersion` to 1.0.

Version Ground Rules

For versioning of the protocol and encoding to be possible, all versions (present and future) of the Ice run time adhere to a few ground rules:

1. Encapsulations always have a six-byte header; the first four bytes are the size of the encapsulation (including the size of the header), followed by two bytes that indicate the major and minor version. How to interpret the remainder of the encapsulation depends on the major and minor version.
2. The first eight bytes of a message header always contain the magic number 'I', 'c', 'e', 'P', followed by four bytes of version information (two bytes for the protocol major and minor number, and two bytes for the encoding major and minor number). How to interpret the remainder of the header and the message body depends on the major and minor version.

These ground rules ensure that all current and future versions of the Ice run time can at least identify the version and size of an encapsulation and a message. This is particularly important for message switches such as IceStorm; by keeping the version and size information in a fixed format, it is possible to forward messages that are, for example, at version 2.0, even though the message switch itself may still be at version 1.0.

See Also

Copyright 2016, ZeroC, Inc.

- Basic Data Encoding
- Protocol Messages
- IceStorm
- IceGrid

Copyright 2016, ZeroC, Inc.

Best Practices

This section provides a number of design guidelines for Ice applications.

Topics

- Optional Values
- Server Implementation Techniques
- Servant Evictors
- Object Life Cycle

Copyright 2016, ZeroC, Inc.

Optional Values

Version 1.1 of the Ice encoding introduces a convenient new feature that allows you to declare data members and parameters as optional. Collectively called *optional values*, this feature provides two main benefits for Ice applications:

- It offers a simple and natively-supported mechanism for representing values that may or may not be present, without consuming any additional space in a message when an optional value is not used.
- It significantly improves the ability to gradually evolve components of an Ice application while maintaining backward compatibility with existing deployments.

On this page:

- Motivation for Optional Values
  - Alternative Strategies for Modeling Optional Values
  - Addressing the Versioning Problem
- Overview of Optional Values
  - Optional Parameters in Operations
  - Optional Data Members in Classes and Exceptions
- Guidelines for using Optional Values
- Backward Compatibility with Encoding Version 1.0
- Using Optional Values Efficiently

Motivation for Optional Values

This section describes the issues that prompted the development of this new feature.

Alternative Strategies for Modeling Optional Values

Prior to Ice 3.5, developers did not have a convenient way to model an optional value. One possibility was to wrap the value in a Slice class:

Slice

```
class OptionalInt
{
    int i;
};

interface Example
{
    void compute(int input, OptionalInt scale);
};
```

The ability to pass a nil value for a class instance means the receiver can interpret nil to mean "not present", and the four bytes required to encode a nil value are a reasonable price to pay. Although this is a workable strategy, it incurs significantly more overhead when the optional value is present because the encoding for class instances is expensive. In the example above, an integer value that would have consumed just four bytes when encoded becomes a class instance that could require ten times as much space in the encoding.

The new "compact" encoding for Slice classes would lower the overhead of this strategy, but it still cannot match the efficiency of the encoding for optional values.

A much more efficient alternative was to use a sequence:

Copyright 2016, ZeroC, Inc.

Slice

```
sequence<int> OptionalInt;

interface Example
{
    void compute(int input, OptionalInt scale);
};
```

Passing an empty sequence signifies the "not present" case and consumes only a single byte in the encoding. When the value is present, passing a sequence with one element adds just one byte of overhead to the encoded value. Despite its minimal overhead, constructing temporary sequences in application code becomes awkward and tedious. Furthermore, depending on the language mapping in use, it is not always obvious that a given sequence parameter actually represents an optional value. Consider the Java language mapping for `compute`:

Java

```
public interface Example : ...
{
    void compute(int input, int[] scale, ...);
}
```

A developer looking at this API for the first time has no indication that the `scale` parameter is optional, which can lead to unintentional misuse and programming errors.

Addressing the Versioning Problem

Aside from the inability to conveniently model an optional value, another motivating factor is versioning. As an application evolves over time, it is only natural for its Slice definitions to change. However, there are only a few kinds of changes you can make to your Slice definitions if your goal is to maintain "on-the-wire" compatibility with existing deployments:

- Adding new interfaces
- Adding new data types that are only used by new operations and interfaces
- Adding new operations to existing interfaces
- Renaming existing parameters and data members
- Renaming structures, sequences, dictionaries, enumerations, and enumerators

The nature of the Ice encoding rules make these changes possible. For example, the encoding does not include the names of data members or parameters, therefore changing their names does not affect backward compatibility. Similarly, the fact that the encoding for an invocation uses the operation's name to select the target operation, and not some other identification scheme such as a numeric index, means you can add new operations to an existing interface without breaking backward compatibility.

Making any of the following modifications is considered a "breaking" change:

- Changing the type of a data member or parameter
- Changing the declaration order of data members or parameters
- Renaming an operation, interface, class, or exception

These changes directly impact the encoding and prevent an existing receiver from successfully decoding the message. Before the introduction of optional values, we recommended using facets to manage a breaking change. Now optional values give you an alternative solution, but only for certain types of changes.

Overview of Optional Values

The optional syntax is the same for both parameters and data members:

`optional(`*`tag`*`)` *`type name`*

The `optional` keyword, new in Ice 3.5, denotes an optional value with the given tag. The value for *`tag`* must be a non-negative integer.

Copyright 2016, ZeroC, Inc.

If your Slice definitions currently use `optional` as an identifier, you can continue using it as an identifier by escaping it as `\optional`.

Data members and parameters not marked as optional are considered *required*; a sender must supply values for all required parameters and members.

Each language mapping defines APIs for passing an optional value, indicating that an optional value is not set, and testing whether an optional value is set. For example, the three strongly-typed languages that Ice currently supports (C++, C#, and Java) all use a similar API: an `Optional` type wraps the optional value and provides methods for examining and changing its value. C++ and C# also define the "marker" value `None` that can be substituted for any optional value to indicate that it is unset. No `Optional` type is necessary for the scripting languages, where you simply pass the value (if set) or the marker value `Unset`. Refer to the language mapping sections for more details on optional values.

A well-behaved program must not assume that an optional data member or parameter always has a value.

An optional value incurs no additional overhead in the encoding when its value is unset. If an optional parameter or member has a value, its overhead depends on several factors, including your choice of tag values.

Optional Parameters in Operations

An operation may declare as optional its return value and any of its parameters. The ordering rules for parameters remain unchanged: all input parameters must precede the output parameters.

The following Slice operation declares a mix of optional and required parameters, as well as an optional return value:

Slice

```
interface Database
{
    optional(3) Record
    lookup(string key, optional(1) Filter f, out optional(2) Cursor matches);
};
```

We can invoke `lookup` in C++ as shown below:

C++

```
DatabasePrx proxy = ...;
Optional<Record> rec;
Optional<Cursor> matches;
rec = proxy->lookup("someKey", IceUtil::None, matches);
if(rec)
{
    // lookup returned a value

    if(matches)
    {
        // there are more matches
    }
}
```

Notice that we pass `IceUtil::None` for the `Filter` value to indicate that this parameter is unset.

Optional Data Members in Classes and Exceptions

Copyright 2016, ZeroC, Inc.

Data members can be declared as optional in classes and exceptions. Structures do not support optional data members.

Consider the following example:

Slice

```
class Account
{
    string accountNo;
    optional(9) Account referrer;
};
```

We can use `Account` in C# as shown below:

C#

```
void printAccount(Account acct)
{
    Console.WriteLine("Account #: " + acct.accountNo);
    if(acct.referrer.HasValue)
        Console.WriteLine("Referrer: " + acct.referrer.Value.accountNo);
    else
        Console.WriteLine("Referrer: None");
}

Account a1 = new Account("100-21807", Ice.Util.None);
printAccount(a1);
Account a2 = new Account("165-75122", a1);
printAccount(a2);
```

Implicit conversion operators in C# allow us to pass a value of the declared type, or `Ice.Util.None` for an unset value, where an optional value is expected.

Guidelines for using Optional Values

Keep the following guidelines in mind while using optional values:

- Consider an optional value as a logical unit composed of its tag and type. The parameter or member name is not included in the Ice encoding, therefore changes to the name do not affect on-the-wire compatibility. Changing the tag of a value is equivalent to adding a new value: the previous iteration of the value still exists as long as there is at least one deployed program that might use it.
- To change the type of an optional value without affecting compatibility with existing deployments, you must also change its tag. You are essentially deprecating the old value and adding a new one.
- Whenever you change the tag of a value, we recommend making a note of the old definition (such as in a Slice comment) so that you do not accidentally reuse its tag.
- Adding a new optional parameter or member does not affect compatibility with existing deployments as long as its tag is not currently in use.
- Removing an optional parameter or member does not affect compatibility with existing deployments if programs are written correctly to test for the presence of the value prior to using it.
- The order of declaration is important for required members and parameters, but is not important for optional members and parameters. If you wish to maintain compatibility with existing deployments, do not change the order of declaration for required members and parameters. Compatibility is not affected by changes to the order of optional members and parameters.

Backward Compatibility with Encoding Version 1.0

Copyright 2016, ZeroC, Inc.

Optional values require version 1.1 of the Ice encoding. If the intended receiver only supports version 1.0 of the encoding, the Ice run time in the sender will omit all optional parameters when marshaling the operation's parameters or results, and omit all optional data members. This behavior makes it possible for you to add optional parameters to an operation, or optional members to a class or exception, without affecting backward compatibility with existing deployments that use version 1.0 of the encoding, while gaining the ability to exchange optional values with new deployments that use version 1.1 of the encoding.

For example, suppose existing deployments use the following class definition:

Slice

```
// Version 1
class Person
{
    string firstName;
    string lastName;
};
```

We can safely add an optional member without affecting compatibility:

Slice

```
// Version 2
class Person
{
    string firstName;
    string lastName;
    optional(1) Date birthDate;
};
```

It is even safe to insert an optional member between two existing members:

Slice

```
// Version 3
class Person
{
    string firstName;
    optional(2) string middleName;
    string lastName;
    optional(1) Date birthDate;
};
```

When sending an instance of `Person` to a receiver that supports only encoding version 1.0, the optional members are omitted and the encoded form matches our initial version of the class.

A similar situation exists for operations. For example, the following three operations are compatible when using encoding version 1.0:

Copyright 2016, ZeroC, Inc.

Slice

```
Person createPerson(string firstName, string lastName); // V1

Person createPerson(string firstName, string lastName, optional(1) Date birthDate);
// V2

Person createPerson(string firstName, optional(2) string middleName, string
lastName,
                    optional(1) Date birthDate); // V3
```

Note however that other types of changes can easily break compatibility.

Using Optional Values Efficiently

The Ice encoding uses a self-describing format to minimize the overhead associated with optional data members and allow for efficient decoding in a receiver. The amount of overhead required for a member depends on the member's type and its tag. The type contributes between one and five bytes of overhead, while the tag value contributes an additional zero to five bytes. Tag values less than 30 add no additional overhead; tag values between 30 and 254 add a single byte, and tag values of 255 or greater add five bytes.

The overhead associated with the member's type is out of your control, but you have direct control of the overhead for tags. To minimize this overhead, use tags in the range zero to 29.

The table below describes the overhead associated with several common Slice types:

Type	Overhead (bytes)
Primitive (including `string`)	1
Proxy	5
Object	1
Sequence of `byte`, `bool`	1
Sequence	2-5
Enumerator	1
Structure	2-5
Dictionary	2-5

Refer to the encoding specification for more details on marshaling optional data members.

See Also

- Data Encoding
- Optional Data Members
- Optional Parameters and Return Values

Copyright 2016, ZeroC, Inc.

Server Implementation Techniques

As mentioned in Servant Locators, instantiating a servant for each Ice object on server start-up is a viable design, provided that you can afford the amount of memory required by the servants, as well as the delay in start-up of the server. However, Ice supports more flexible mappings between Ice objects and servants; these alternate mappings allow you to precisely control the trade-off between memory consumption, scalability, and performance. We outline a few of the more common implementation techniques here.

On this page:

- Incremental Server Initialization
- Implementing a Server using Default Servants
  - Overriding ice_ping
- Combining Server Implementation Techniques

Incremental Server Initialization

If you use a servant locator, the servant returned by `locate` is used only for the current request, that is, the Ice run time does not add the servant to the Active Servant Map (ASM). Of course, this means that if another request comes in for the same Ice object, `locate` must again retrieve the object state and instantiate a servant. A common implementation technique is to add each servant to the ASM as part of `locate`. This means that only the first request for each Ice object triggers a call to `locate`; thereafter, the servant for the corresponding Ice object can be found in the ASM and the Ice run time can immediately dispatch another incoming request for the same Ice object without having to call the servant locator.

An implementation of `locate` to do this would look something like the following:

C++

```
Ice::ObjectPtr
MyServantLocator::locate(const Ice::Current& c, Ice::LocalObjectPtr&)
{
    // Get the object identity. (We use the name member
    // as the database key.)
    //
    std::string name = c.id.name;

    // Use the identity to retrieve the state from the database.
    //
    ServantDetails d;
    try {
        d = DB_lookup(name);
    } catch (const DB_error&)
        return 0;
    }

    // We have the state, instantiate a servant.
    //
    Ice::ObjectPtr servant = new PhoneEntryI(d);

    // Add the servant to the ASM.
    //
    c.adapter->add(servant, c.id);      // NOTE: Incorrect!

    return servant;
}
```

This is almost identical to the implementation seen in our earlier example — the only difference is that we also add the servant to the ASM by calling `ObjectAdapter::add`. Unfortunately, this implementation is wrong because it suffers from a race condition. Consider the situation where we do not have a servant for a particular Ice object in the ASM, and two clients more or less simultaneously send a request

Copyright 2016, ZeroC, Inc.

for the same Ice object. It is entirely possible for the thread scheduler to schedule the two incoming requests such that the Ice run time completes the lookup in the ASM for both requests and, for each request, concludes that no servant is in memory. The net effect is that `locate` will be called twice for the same Ice object, and our servant locator will instantiate two servants instead of a single servant. Because the second call to `ObjectAdapter::add` will raise an `AlreadyRegisteredException`, only one of the two servants will be added to the ASM.

Of course, this is hardly the behavior we expect. To avoid the race condition, our implementation of `locate` must check whether a concurrent invocation has already instantiated a servant for the incoming request and, if so, return that servant instead of instantiating a new one. The Ice run time provides the `ObjectAdapter::find` operation to allow us to test whether an entry for a specific identity already exists in the ASM:

Slice

```
module Ice {
    local interface ObjectAdapter {
        // ...

        Object find(Identity id);

        // ...
    };
};
```

`find` returns the servant if it exists in the ASM and null, otherwise. Using this lookup function, together with a mutex, allows us to correctly implement `locate`. The class definition of our servant locator now has a private mutex so we can establish a critical region inside `locate`:

C++

```
class MyServantLocator : public virtual Ice::ServantLocator {
public:

    virtual Ice::ObjectPtr locate(const Ice::Current& c, Ice::LocalObjectPtr&);

    // Declaration of finished() and deactivate() here...

private:
    IceUtil::Mutex _m;
};
```

The `locate` member locks the mutex and tests whether a servant is already in the ASM: if so, it returns that servant; otherwise, it instantiates a new servant and adds it to the ASM as before:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::ObjectPtr
MyServantLocator::locate(const Ice::Current& c, Ice::LocalObjectPtr&)
{
    IceUtil::Mutex::Lock lock(_m);

    // Check if we have instantiated a servant already.
    //
    Ice::ObjectPtr servant = c.adapter.find(c.id);

    if (!servant) {     // We don't have a servant already

        // Instantiate a servant.
        //
        ServantDetails d;
        try {
           d = DB_lookup(c.id.name);
        } catch (const DB_error&) {
           return 0;
        }
        servant = new PhoneEntryI(d);

        // Add the servant to the ASM.
        //
        c.adapter->add(servant, c.id);
    }

    return servant;
}
```

The Java version of this locator is almost identical, but we use the `synchronized` qualifier instead of a mutex to make `locate` a critical region:

Copyright 2016, ZeroC, Inc.

Java

```
synchronized public Ice.Object
locate(Ice.Current c, Ice.LocalObjectHolder cookie)
{
    // Check if we have instantiated a servant already.
    //
    Ice.Object servant = c.adapter.find(c.id);

    if (servant == null) { // We don't have a servant already

        // Instantiate a servant
        //
        ServantDetails d;
        try {
            d = DB.lookup(c.id.name);
        } catch (DB.error&) {
            return null;
        }
        servant = new PhoneEntryI(d);

        // Add the servant to the ASM.
        //
        c.adapter.add(servant, c.id);
    }

    return servant;
}
```

In C#, you can place the body of `locate` into a `lock(this)` statement.

Using a servant locator that adds the servant to the ASM has a number of advantages:

- Servants are instantiated on demand, so the cost of initializing the servants is spread out over many invocations instead of being incurred all at once during server start-up.
- The memory requirements for the server are reduced because servants are instantiated only for those Ice objects that are actually accessed by clients. If clients only access a subset of the total number of Ice objects, the memory savings can be substantial.

In general, incremental initialization is beneficial if instantiating servants during start-up is too slow. The memory savings can be worthwhile as well but, as a rule, are realized only for comparatively short-lived servers: for long-running servers, chances are that, sooner or later, every Ice object will be accessed by some client or another; in that case, there are no memory savings because we end up with an instantiated servant for every Ice object regardless.

Implementing a Server using Default Servants

Default servants are a very effective tool for conserving memory when a server hosts a large number of Ice objects.

To create a default servant implementation, we need as many default servants as there are non-abstract interfaces in the system. For example, for our file system application, we require two default servants, one for directories and one for files. In addition, the object identities we create use the `category` member of the object identity to encode the type of interface of the corresponding Ice object. The value of the category field can be anything that identifies the interface, such as the '`d`' and '`f`' convention we suggested earlier. Alternatively, you could use `"Directory"` and `"File"`, or use the type ID of the corresponding interface, such as `"::Filesystem::Directory"` and `"::File system::File"`. The `name` member of the object identity must be set to whatever identifier we can use to retrieve the persistent state of each directory and file from secondary storage. (For our file system application, we used a UUID as a unique identifier.)

Registration of the default servants is as follows:

Copyright 2016, ZeroC, Inc.

C++

```
adapter->addDefaultServant(new DirectoryI, "d");
adapter->addDefaultServant(new FileI, "f");
```

All the action happens in the implementation of the operations, using the following steps for each operation:

1. Use the passed `Current` object to get the identity for the current request.
2. Use the `name` member of the identity to locate the persistent state of the servant on secondary storage. If no record can be found for the identity, throw an `ObjectNotExistException`.
3. Implement the operation to operate on that retrieved state (returning the state or updating the state as appropriate for the operation).

This might look something like the following:

C++

```
Filesystem::NodeSeq
Filesystem::DirectoryI::list(const Ice::Current& c) const
{
    // Use the identity of the directory to retrieve
    // its contents.
    DirectoryContents dc;
    try {
        dc = DB_getDirectory(c.id.name);
    } catch(const DB_error&) {
        throw Ice::ObjectNotExistException(__FILE__, __LINE__);
    }

    // Use the records retrieved from the database to
    // initialize return value.
    //
    FileSystem::NodeSeq ns;
    // ...

    return ns;
}
```

Note that the servant implementation is completely stateless: the only state it operates on is the identity of the Ice object for the current request (and that identity is passed as part of the `Current` parameter).

Overriding `ice_ping`

We recommended that a default servant implementation take steps to preserve the semantics of the `ice_ping` operation, which is used to test whether an Ice object exists. If a default servant fails to override `ice_ping`, clients may mistakenly believe that a non-existent Ice object still exists. The code below demonstrates how we can override the operation in our file system application:

Copyright 2016, ZeroC, Inc.

C++

```
void
Filesystem::DirectoryI::ice_ping(const Ice::Current& c) const
{
    try {
       d = DB_lookup(c.id.name);
    } catch (const DB_error&) {
       throw Ice::ObjectNotExistException(__FILE__, __LINE__);
    }
}
```

It is good practice to override `ice_ping` if you are using default servants.

Combining Server Implementation Techniques

Depending on the nature of your application, you may be able to steer a middle path that provides better performance while keeping memory requirements low: if your application has a number of frequently-accessed objects that are performance-critical, you can add servants for those objects to the ASM. If you store the state of these objects in data members inside the servants, you effectively have a cache of these objects.

The remaining, less-frequently accessed objects can be implemented with a default servant. For example, in our file system implementation, we could choose to instantiate directory servants permanently, but to have file objects implemented with a default servant. This provides efficient navigation through the directory tree and incurs slower performance only for the (presumably less frequent) file accesses.

This technique could be augmented with a cache of recently-accessed files, along similar lines to the buffer pool used by the Unix kernel [1]. The point is that you can combine use of the ASM with servant locators and default servants to precisely control the trade-offs among scalability, memory consumption, and performance to suit the needs of your application.

See Also

- Servant Locators
- Servant Locator Example
- Servant Activation and Deactivation
- Using Identity Categories with Servant Locators
- Default Servants

References

1. McKusick, M. K., et al. 1996. The Design and Implementation of the 4.4BSD Operating System. Reading, MA: Addison-Wesley.

Copyright 2016, ZeroC, Inc.

Servant Evictors

A particularly interesting use of a servant locator is as an *evictor* [1]. An evictor is a servant locator that maintains a cache of servants:

- Whenever a request arrives (that is, `locate` is called by the Ice run time), the evictor checks to see whether it can find a servant for the request in its cache. If so, it returns the servant that is already instantiated in the cache; otherwise, it instantiates a servant and adds it to the cache.
- The cache is a queue that is maintained in least-recently used (LRU) order: the least-recently used servant is at the tail of the queue, and the most-recently used servant is at the head of the queue. Whenever a servant is returned from or added to the cache, it is moved from its current queue position to the head of the queue, that is, the "newest" servant is always at the head, and the "oldest" servant is always at the tail.
- The queue has a configurable length that corresponds to how many servants will be held in the cache; if a request arrives for an Ice object that does not have a servant in memory and the cache is full, the evictor removes the least-recently used servant at the tail of the queue from the cache in order to make room for the servant about to be instantiated at the head of the queue.

The figure below illustrates an evictor with a cache size of five after five invocations have been made, for object identities 1 to 5, in that order.

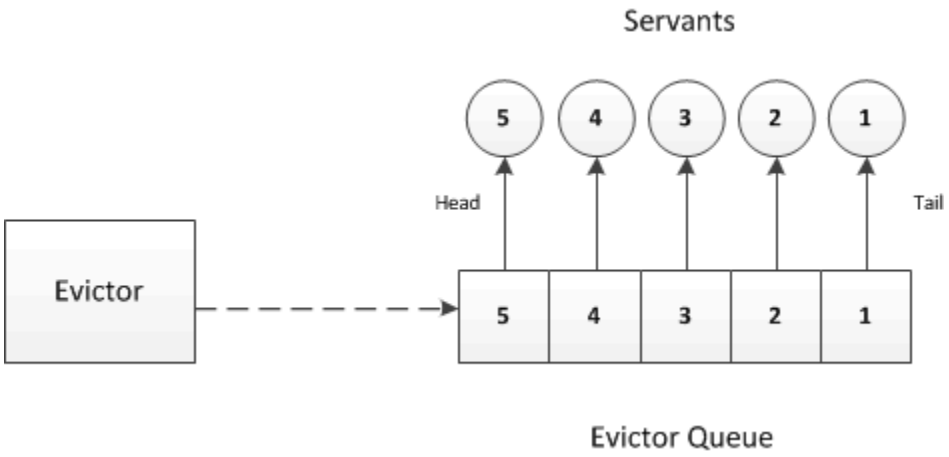

An evictor after five invocations for object identities 1 to 5.

At this point, the evictor has instantiated five servants, and has placed each servant onto the evictor queue. Because requests were sent by the client for object identities 1 to 5 (in that order), servant 5 ends up at the head of the queue (at the most-recently used position), and servant 1 ends up at the tail of the queue (at the least-recently used position).

Assume that the client now sends a request for servant 3. In this case, the servant is found on the evictor queue and moved to the head position. The resulting ordering is shown below:

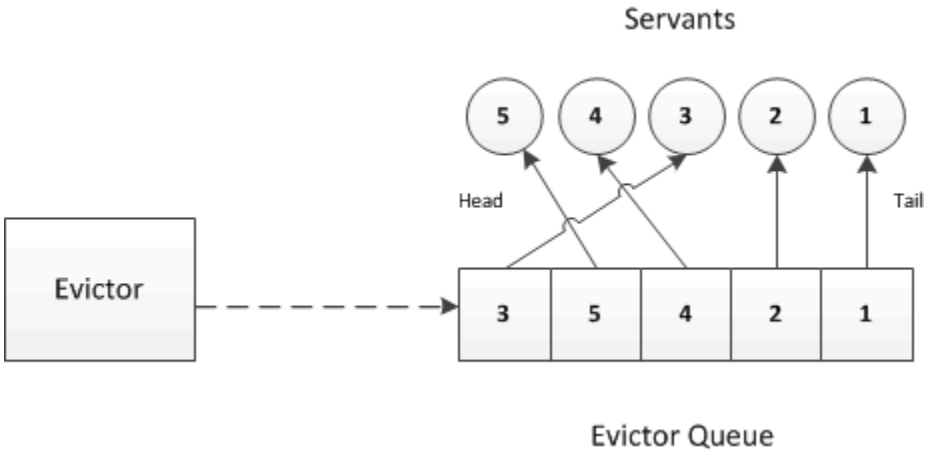

The evictor after accessing servant 3.

Assume that the next client request is for object identity 6. The evictor queue is fully populated, so the evictor creates a servant for object identity 6, places that servant at the head of the queue, and evicts the servant with identity 1 (the least-recently used servant) at the tail of the queue, as you can see here:

Copyright 2016, ZeroC, Inc.

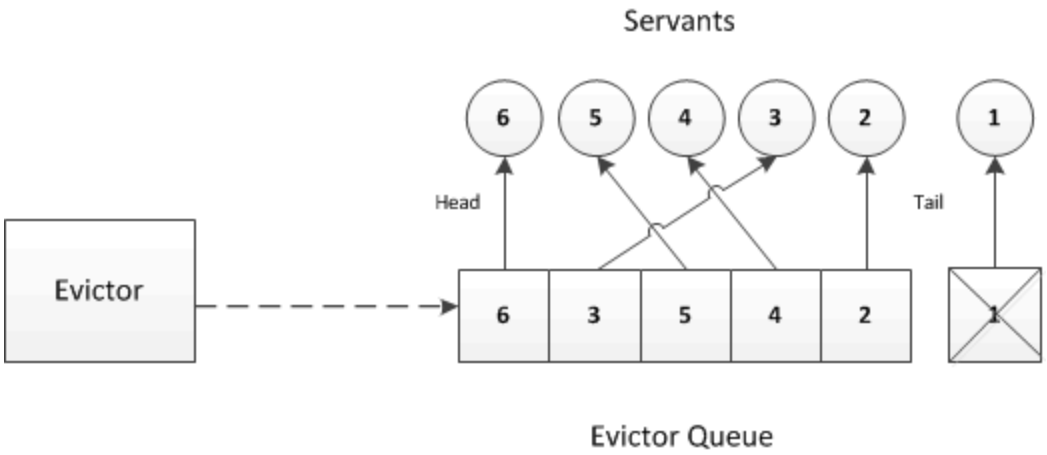

The evictor after evicting servant 1.

The evictor pattern combines the advantages of the ASM with the advantages of a default servant: provided that the cache size is sufficient to hold the working set of servants in memory, most requests are served by an already instantiated servant, without incurring the overhead of creating a servant and accessing the database to initialize servant state. By setting the cache size, you can control the trade-off between performance and memory consumption as appropriate for your application.

The following pages show how to implement an evictor in several languages. (You can also find the source code for the evictor with the code examples for this manual in the Ice distribution.)

Topics

- Implementing a Servant Evictor in C++
- Implementing a Servant Evictor in Java
- Implementing a Servant Evictor in C-Sharp

See Also

- Servant Locators
- Default Servants

References

1. Henning, M., and S. Vinoski. 1999. Advanced CORBA Programming with C++. Reading, MA: Addison-Wesley.

Copyright 2016, ZeroC, Inc.

Implementing a Servant Evictor in C++

On this page:

- The EvictorBase Class in C++
- Using Servant Evictors in C++

The `EvictorBase` Class in C++

The evictor we show here is designed as an abstract base class: in order to use it, you derive a class from the `EvictorBase` base class and implement two methods that are called by the evictor when it needs to add or evict a servant. This leads to a class definition as follows:

C++

```
class EvictorBase : public Ice::ServantLocator {
public:
    EvictorBase(int size = 1000);

    virtual Ice::ObjectPtr locate(const Ice::Current& c,
Ice::LocalObjectPtr& cookie);

    virtual void finished(const Ice::Current& c, const Ice::ObjectPtr&,
                          const Ice::LocalObjectPtr& cookie);

    virtual void deactivate(const std::string&);

protected:
    virtual Ice::ObjectPtr add(const Ice::Current&, Ice::LocalObjectPtr&) = 0;

    virtual void evict(const Ice::ObjectPtr&, const Ice::LocalObjectPtr&) = 0;

private:
    // ...
};

typedef IceUtil::Handle<EvictorBase> EvictorBasePtr;
```

Note that the evictor has a constructor that sets the size of the queue, with a default argument to set the size to 1000.

The `locate`, `finished`, and `deactivate` functions are inherited from the `ServantLocator` base class; these functions implement the logic to maintain the queue in LRU order and to add and evict servants as needed.

The `add` and `evict` functions are called by the evictor when it needs to add a new servant to the queue and when it evicts a servant from the queue. Note that these functions are pure virtual, so they must be implemented in a derived class. The job of `add` is to instantiate and initialize a servant for use by the evictor. The `evict` function is called by the evictor when it evicts a servant, allowing the subclass to perform any cleanup. Note that `add` can return a cookie that the evictor passes to `evict`, so you can move context information from `add` to `evict`.

Next, we need to consider the data structures that are needed to support our evictor implementation. We require two main data structures:

1. A map that maps object identities to servants, so we can efficiently decide whether we have a servant for an incoming request in memory or not.
2. A list that implements the evictor queue. The list is kept in LRU order at all times.

The evictor map not only stores servants but also keeps track of some administrative information:

1. The map stores the cookie that is returned from `add`, so we can pass that same cookie to `evict`.
2. The map stores an iterator into the evictor queue that marks the position of the servant in the queue. Storing the queue position is not strictly necessary — we store the position for efficiency reasons because it allows us to locate a servant's position in the queue in constant time instead of having to search through the queue in order to maintain its LRU property.

Copyright 2016, ZeroC, Inc.

3. The map stores a use count that is incremented whenever an operation is dispatched into a servant, and decremented whenever an operation completes.

The need for the use count deserves some extra explanation: suppose a client invokes a long-running operation on an Ice object with identity *I*. In response, the evictor adds a servant for *I* to the evictor queue. While the original invocation is still executing, other clients invoke operations on various Ice objects, which leads to more servants for other object identities being added to the queue. As a result, the servant for identity *I* gradually migrates toward the tail of the queue. If enough client requests for other Ice objects arrive while the operation on object *I* is still executing, the servant for *I* could be evicted while it is still executing the original request.

By itself, this will not do any harm. However, if the servant is evicted and a client then invokes another request on object *I*, the evictor would have no idea that a servant for *I* is still around and would add a second servant for *I*. However, having two servants for the same Ice object in memory is likely to cause problems, especially if the servant's operation implementations write to a database.

The use count allows us to avoid this problem: we keep track of how many requests are currently executing inside each servant and, while a servant is busy, avoid evicting that servant. As a result, the queue size is not a hard upper limit: long-running operations can temporarily cause more servants than the limit to appear in the queue. However, as soon as excess servants become idle, they are evicted as usual.

The evictor queue does not store the identity of the servant. Instead, the entries on the queue are iterators into the evictor map. This is useful when the time comes to evict a servant: instead of having to search the map for the identity of the servant to be evicted, we can simply delete the map entry that is pointed at by the iterator at the tail of the queue. We can get away with storing an iterator into the evictor queue as part of the map, and storing an iterator into the evictor map as part of the queue because both `std::list` and `std::map` do not invalidate forward iterators when we add or delete entries (except for invalidating iterators that point at a deleted entry, of course).

Reverse iterators *can* be invalidated by modification of list entries: if a reverse iterator points at `rend` and the element at the head of the list is erased, the iterator pointing at `rend` is invalidated.

Finally, our `locate` and `finished` implementations will need to exchange a cookie that contains a smart pointer to the entry in the evictor map. This is necessary so that `finished` can decrement the servant's use count.

This leads to the following definitions in the private section of our evictor:

Copyright 2016, ZeroC, Inc.

C++

```
class EvictorBase : public Ice::ServantLocator {
    // ...

private:

    struct EvictorEntry;
    typedef IceUtil::Handle<EvictorEntry> EvictorEntryPtr;

    typedef std::map<Ice::Identity, EvictorEntryPtr> EvictorMap;
    typedef std::list<EvictorMap::iterator> EvictorQueue;

    struct EvictorEntry : public Ice::LocalObject
    {
        Ice::ObjectPtr servant;
        Ice::LocalObjectPtr userCookie;
        EvictorQueue::iterator queuePos;
        int useCount;
    };

    EvictorMap _map;
    EvictorQueue _queue;
    Ice::Int _size;

    IceUtil::Mutex _mutex;

    void evictServants();
};
```

Note that the evictor stores the evictor map, queue, and the queue size in the private data members `_map`, `_queue`, and `_size`. In addition, we use a private `_mutex` data member so we can correctly serialize access to the evictor's data structures.

The `evictServants` member function takes care of evicting servants when the queue length exceeds its limit — we will discuss this function in more detail shortly.

The `EvictorEntry` structure serves as the cookie that we pass from `locate` to `finished`; it stores the servant, the servant's position in the evictor queue, the servant's use count, and the cookie that we pass from `add` to `evict`.

The implementation of the constructor is trivial. The only point of note is that we ignore negative sizes:

C++

```
EvictorBase::EvictorBase(Ice::Int size) : _size(size)
{
    if (_size < 0)
        _size = 1000;
}
```

We could have stored the size as a `size_t` instead. However, for consistency with the Java implementation, which cannot use unsigned integers, we use `Ice::Int` to store the size.

Copyright 2016, ZeroC, Inc.

Almost all the action of the evictor takes place in the implementation of `locate`:

C++

```
Ice::ObjectPtr
EvictorBase::locate(const Ice::Current& c, Ice::LocalObjectPtr& cookie)
{
    IceUtil::Mutex::Lock lock(_mutex);

    //
    // Check if we have a servant in the map already.
    //
    EvictorEntryPtr entry;
    EvictorMap::iterator i = _map.find(c.id);
    if (i != _map.end()) {
        //
        // Got an entry already, dequeue the entry from
        // its current position.
        //
        entry = i->second;
        _queue.erase(entry->queuePos);
    } else {
        //
        // We do not have an entry. Ask the derived class to
        // instantiate a servant and add a new entry to the map.
        //
        entry = new EvictorEntry;
        entry->servant = add(c, entry->userCookie); // Down-call
        if (!entry->servant) {
            return 0;
        }
        entry->useCount = 0;
        i = _map.insert(std::make_pair(c.id, entry)).first;
    }

    //
    // Increment the use count of the servant and enqueue
    // the entry at the front, so we get LRU order.
    //
    ++(entry->useCount);
    entry->queuePos = _queue.insert(_queue.begin(), i);

    cookie = entry;

    return entry->servant;
}
```

The first step in `locate` is to lock the `_mutex` data member. This protects the evictor's data structures from concurrent access. The next step is to instantiate a smart pointer to an `EvictorEntry`. That smart pointer acts as the cookie that is returned from `locate` and will be passed by the Ice run time to the corresponding call to `finished`. That same smart pointer is also the value type of our map entries, so we do not store two copies of the same information redundantly — instead, smart pointers ensure that a single copy of each `EvictorEntry` str ucture is shared by both the cookie and the map.

Copyright 2016, ZeroC, Inc.

The next step is to look in the evictor map to see whether we already have an entry for this object identity. If so, we remove the entry from its current queue position.

Otherwise, we do not have an entry for this object identity yet, so we have to create one. The code creates a new evictor entry, and then calls `add` to get a new servant. This is a down-call to the concrete class that will be derived from `EvictorBase`. The implementation of `add` must attempt to locate the object state for the Ice object with the identity passed inside the `Current` object and either return a servant as usual, or return null or throw an exception to indicate failure. If `add` returns null, we return zero to let the Ice run time know that no servant could be found for the current request. If `add` succeeds, we initialize the entry's use count to zero and insert the entry into the evictor map.

The last few lines of `locate` add the entry for the current request to the head of the evictor queue to maintain its LRU property, increment the use count of the entry, set the cookie that is returned from locate to point at the `EvictorEntry`, and finally return the servant to the Ice run time.

The implementation of `finished` is comparatively simple. It decrements the use count of the entry and then calls `evictServants` to get rid of any servants that might need to be evicted:

C++

```
void
EvictorBase::finished(const Ice::Current&,
                      const Ice::ObjectPtr&,
                      const Ice::LocalObjectPtr& cookie)
{
    IceUtil::Mutex::Lock lock(_mutex);

    EvictorCookiePtr ec = EvictorCookiePtr::dynamicCast(cookie);

    // Decrement use count and check if
    // there is something to evict.
    //
    --(ec->entry->useCount);
    evictServants();
}
```

In turn, `evictServants` examines the evictor queue: if the queue length exceeds the evictor's size, the excess entries are scanned. Any entries with a zero use count are then evicted:

Copyright 2016, ZeroC, Inc.

C++

```
void
EvictorBase::evictServants()
{
    //
    // If the evictor queue has grown larger than the limit,
    // look at the excess elements to see whether any of them
    // can be evicted.
    //
    EvictorQueue::reverse_iterator p = _queue.rbegin();
    int excessEntries = static_cast<int>(_map.size() - _size);

    for (int i = 0; i < excessEntries; ++i) {
        EvictorMap::iterator mapPos = *p;
        if (mapPos->second->useCount == 0) {
            evict(mapPos->second->servant, mapPos->second->userCookie);
            p = EvictorQueue::reverse_iterator(_queue.erase(mapPos->second->queuePo
s));
            _map.erase(mapPos);
        }
        else
            ++p;
    }
}
```

The code scans the excess entries, starting at the tail of the evictor queue. If an entry has a zero use count, it is evicted: after calling the `evict` member function in the derived class, the code removes the evicted entry from both the map and the queue.

Finally, the implementation of `deactivate` sets the evictor size to zero and then calls `evictServants`. This results in eviction of all servants. The Ice run time guarantees to call `deactivate` only once no more requests are executing in an object adapter; as a result, it is guaranteed that all entries in the evictor will be idle and therefore will be evicted.

C++

```
void
EvictorBase::deactivate(const std::string& category)
{
    IceUtil::Mutex::Lock lock(_mutex);

    _size = 0;
    evictServants();
}
```

Note that, with this implementation of `evictServants`, we only scan the tail section of the evictor queue for servants to evict. If we have long-running operations, this allows the number of servants in the queue to remain above the evictor size if the servants in the tail section have a non-zero use count. This means that, even immediately after calling `evictServants`, the queue length can still exceed the evictor size.

We can adopt a more aggressive strategy for eviction: instead of scanning only the excess entries in the queue, if, after looking in the tail section of the queue, we still have more servants in the queue than the queue size, we keep scanning for servants with a zero use count

Copyright 2016, ZeroC, Inc.

until the queue size drops below the limit. This alternative version of `evictServants` looks as follows:

C++

```
void
EvictorBase::evictServants()
{
    //
    // If the evictor queue has grown larger than the limit,
    // try to evict servants until the length drops
    // below the limit.
    //
    EvictorQueue::reverse_iterator p = _queue.rbegin();
    int numEntries = static_cast<int>_map.size();

    for (int i = 0; i < numEntries && _map.size() > _size; ++i) {
        EvictorMap::iterator mapPos = *p;
        if (mapPos->second->useCount == 0) {
            evict(mapPos->second->servant, mapPos->second->userCookie);
            p = EvictorQueue::reverse_iterator(_queue.erase(mapPos->second->queuePo
s));
            _map.erase(mapPos);
        }
        else
            ++p;
    }
}
```

The only difference in this version is that the terminating condition for the `for`-loop has changed: instead of scanning only the excess entries for servants with a use count, this version keeps scanning until the evictor size drops below the limit.

Which version is more appropriate depends on your application: if locating and evicting servants is expensive, and memory is not at a premium, the first version (which only scans the tail section) is more appropriate; if you want to keep memory consumption to a minimum, the second version is more appropriate. Also keep in mind that the difference between the two versions is significant only if you have long-running operations and many concurrent invocations from clients; otherwise, there is no point in more aggressively scanning for servants to remove because they are going to become idle again very quickly and get evicted as soon as the next request arrives.

Using Servant Evictors in C++

Using a servant evictor is simply a matter of deriving a class from `EvictorBase` and implementing the `add` and `evict` methods. You can turn a servant locator into an evictor by simply taking the code that you wrote for `locate` and placing it into `add` — `EvictorBase` then takes care of maintaining the cache in least-recently used order and evicting servants as necessary. Unless you have clean-up requirements for your servants (such as closing network connections or database handles), the implementation of `evict` can be left empty.

One of the nice aspects of evictors is that you do not need to change anything in your servant implementation: the servants are ignorant of the fact that an evictor is in use. This makes it very easy to add an evictor to an already existing code base with little disturbance of the source code.

Evictors can provide substantial performance improvements over default servants: especially if initialization of servants is expensive (for example, because servant state must be initialized by reading from a network), an evictor performs much better than a default servant, while keeping memory requirements low.

See Also

- Servant Evictors
- Object Identity
- Default Servants

Copyright 2016, ZeroC, Inc.

Implementing a Servant Evictor in Java

On this page:

- The EvictorBase Class in Java
- Using Servant Evictors in Java

The `EvictorBase` Class in Java

The evictor we show here is designed as an abstract base class: in order to use it, you derive a class from the `Evictor.EvictorBase` base class and implement two methods that are called by the evictor when it needs to add or evict a servant. This leads to a class definition as follows:

Copyright 2016, ZeroC, Inc.

Java

```
package Evictor;

public abstract class EvictorBase implements Ice.ServantLocator
{
    public
    EvictorBase()
    {
        _size = 1000;
    }

    public
    EvictorBase(int size)
    {
        _size = size < 0 ? 1000 : size;
    }

    public abstract Ice.Object
    add(Ice.Current c, Ice.LocalObjectHolder cookie);

    public abstract void
    evict(Ice.Object servant, java.lang.Object cookie);

    synchronized public final Ice.Object
    locate(Ice.Current c, Ice.LocalObjectHolder cookie)
    {
        // ...
    }

    synchronized public final void
    finished(Ice.Current c, Ice.Object o, java.lang.Object cookie)
    {
        // ...
    }

    synchronized public final void
    deactivate(String category)
    {
        // ...
    }

    // ...

    private int _size;
}
```

Note that the evictor has constructors to set the size of the queue, with a default size of 1000.

The `locate`, `finished`, and `deactivate` methods are inherited from the `ServantLocator` base class; these methods implement the logic to maintain the queue in LRU order and to add and evict servants as needed. The methods are synchronized, so the evictor's internal data structures are protected from concurrent access.

Copyright 2016, ZeroC, Inc.

The `add` and `evict` methods are called by the evictor when it needs to add a new servant to the queue and when it evicts a servant from the queue. Note that these functions are abstract, so they must be implemented in a derived class. The job of `add` is to instantiate and initialize a servant for use by the evictor. The `evict` function is called by the evictor when it evicts a servant, allowing the subclass to perform any cleanup. Note that `add` can return a cookie that the evictor passes to `evict`, so you can move context information from `add` to `evict`.

Next, we need to consider the data structures that are needed to support our evictor implementation. We require two main data structures:

1. A map that maps object identities to servants, so we can efficiently decide whether we have a servant for an incoming request in memory or not.
2. A list that implements the evictor queue. The list is kept in LRU order at all times.

The evictor map not only stores servants but also keeps track of some administrative information:

1. The map stores the cookie that is returned from `add`, so we can pass that same cookie to `evict`.
2. The map stores an iterator into the evictor queue that marks the position of the servant in the queue.
3. The map stores a use count that is incremented whenever an operation is dispatched into a servant, and decremented whenever an operation completes.

The last two points deserve some extra explanation.

- The evictor queue must be maintained in least-recently used order, that is, every time an invocation arrives and we find an entry for the identity in the evictor map, we also must locate the servant's identity on the evictor queue and move it to the front of the queue. However, scanning for that entry is inefficient because it requires *O*(*n*) time. To get around this, we store an iterator in the evictor map that marks the corresponding entry's position in the evictor queue. This allows us to dequeue the entry from its current position and enqueue it at the head of the queue in *O*(1) time.

  Unfortunately, the various lists provided by `java.util` do not allow us to keep an iterator to a list position without invalidating that iterator as the list is updated. To deal with this, we use a special-purpose linked list implementation, `Evictor.LinkedList`, that does not have this limitation. `LinkedList` has an interface similar to `java.util.LinkedList` but does not invalidate iterators other than iterators that point at an element that is removed. For brevity, we do not show the implementation of this list here — you can find the implementation in the code examples for this manual in the Ice distribution.

- We maintain a use count as part of the map in order to avoid incorrect eviction of servants. Suppose a client invokes a long-running operation on an Ice object with identity *I*. In response, the evictor adds a servant for *I* to the evictor queue. While the original invocation is still executing, other clients invoke operations on various Ice objects, which leads to more servants for other object identities being added to the queue. As a result, the servant for identity *I* gradually migrates toward the tail of the queue. If enough client requests for other Ice objects arrive while the operation on object *I* is still executing, the servant for *I* could be evicted while it is still executing the original request.

  By itself, this will not do any harm. However, if the servant is evicted and a client then invokes another request on object *I*, the evictor would have no idea that a servant for *I* is still around and would add a second servant for *I*. However, having two servants for the same Ice object in memory is likely to cause problems, especially if the servant's operation implementations write to a database.

  The use count allows us to avoid this problem: we keep track of how many requests are currently executing inside each servant and, while a servant is busy, avoid evicting that servant. As a result, the queue size is not a hard upper limit: long-running operations can temporarily cause more servants than the limit to appear in the queue. However, as soon as excess servants become idle, they are evicted as usual.

Finally, our `locate` and `finished` implementations will need to exchange a cookie that contains a smart pointer to the entry in the evictor map. This is necessary so that `finished` can decrement the servant's use count.

This leads to the following definitions in the private section of our evictor:

Copyright 2016, ZeroC, Inc.

Java

```
package Evictor;

public abstract class EvictorBase implements Ice.ServantLocator
{
    // ...

    private class EvictorEntry
    {
        Ice.Object servant;
        java.lang.Object userCookie;
        java.util.Iterator<Ice.Identity> queuePos;
        int useCount;
    }

    private void evictServants()
    {
        // ...
    }

    private java.util.Map<Ice.Identity, EvictorEntry> _map =
        new java.util.HashMap<Ice.Identity, EvictorEntry>();
    private Evictor.LinkedList<Ice.Identity> _queue =
        new Evictor.LinkedList<Ice.Identity>();
    private int _size;
}
```

Note that the evictor stores the evictor map, queue, and the queue size in the private data members `_map`, `_queue`, and `_size`. The map key is the identity of the Ice object, and the lookup value is of type `EvictorEntry`. The queue simply stores identities, of type `Ice::Identity`.

The `evictServants` member function takes care of evicting servants when the queue length exceeds its limit — we will discuss this function in more detail shortly.

Almost all the action of the evictor takes place in the implementation of `locate`:

Copyright 2016, ZeroC, Inc.

Java

```
synchronized public final Ice.Object
locate(Ice.Current c, Ice.LocalObjectHolder cookie)
{
    //
    // Check if we have a servant in the map already.
    //
    EvictorEntry entry = _map.get(c.id);
    if (entry != null) {
        //
        // Got an entry already, dequeue the entry from
        // its current position.
        //
        entry.queuePos.remove();
    } else {
        //
        // We do not have entry. Ask the derived class to
        // instantiate a servant and add a new entry to the map.
        //
        entry = new EvictorEntry();
        Ice.LocalObjectHolder cookieHolder = new Ice.LocalObjectHolder();
        entry.servant = add(c, cookieHolder); // Down-call
        if (entry.servant == null) {
            return null;
        }
        entry.userCookie = cookieHolder.value;
        entry.useCount = 0;
        _map.put(c.id, entry);
    }

    //
    // Increment the use count of the servant and enqueue
    // the entry at the front, so we get LRU order.
    //
    ++(entry.useCount);
    _queue.addFirst(c.id);
    entry.queuePos = _queue.iterator();
    entry.queuePos.next(); // Position iterator on the element.

    cookie.value = entry;
    return entry.servant;
}
```

The code uses an `EvictorEntry` as the cookie that is returned from `locate` and will be passed by the Ice run time to the corresponding call to `finished`.

We first look for an existing entry in the evictor map, using the object identity as the key. If we have an entry in the map already, we dequeue the corresponding identity from the evictor queue. (The `queuePos` member of `EvictorEntry` is an iterator that marks that entry's position in the evictor queue.)

Otherwise, we do not have an entry in the map, so we create a new one and call the `add` method. This is a down-call to the concrete class that will be derived from `EvictorBase`. The implementation of `add` must attempt to locate the object state for the Ice object with the identity passed inside the `Current` object and either return a servant as usual, or return null or throw an exception to indicate failure. If `add` returns

Copyright 2016, ZeroC, Inc.

null, we return null to let the Ice run time know that no servant could be found for the current request. If `add` succeeds, we initialize the entry's use count to zero and insert the entry into the evictor map.

The final few lines of code increment the entry's use count, add the entry at the head of the evictor queue, store the entry's position in the queue, and assign the entry to the cookie that is returned from `locate`, before returning the servant to the Ice run time.

The implementation of `finished` is comparatively simple. It decrements the use count of the entry and then calls `evictServants` to get rid of any servants that might need to be evicted:

Java

```
synchronized public final void
finished(Ice.Current c, Ice.Object o, java.lang.Object cookie)
{
    EvictorEntry entry = (EvictorEntry)cookie;

    // Decrement use count and check if
    // there is something to evict.
    //
    --(entry.useCount);
    evictServants();
}
```

In turn, `evictServants` examines the evictor queue: if the queue length exceeds the evictor's size, the excess entries are scanned. Any entries with a zero use count are then evicted:

Java

```
private void evictServants()
{
    //
    // If the evictor queue has grown larger than the limit,
    // look at the excess elements to see whether any of them
    // can be evicted.
    //
    java.util.Iterator<Ice.Identity> p = _queue.riterator();
    int excessEntries = _map.size() - _size;
    for (int i = 0; i < excessEntries; ++i) {
        Ice.Identity id = p.next();
        EvictorEntry e = _map.get(id);
        if (e.useCount == 0) {
            evict(e.servant, e.userCookie); // Down-call
            e.queuePos.remove();
            _map.remove(id);
        }
    }
}
```

The code scans the excess entries, starting at the tail of the evictor queue. If an entry has a zero use count, it is evicted: after calling the `evict` member function in the derived class, the code removes the evicted entry from both the map and the queue.

Finally, the implementation of `deactivate` sets the evictor size to zero and then calls `evictServants`. This results in eviction of all servants. The Ice run time guarantees to call `deactivate` only once no more requests are executing in an object adapter; as a result, it is

Copyright 2016, ZeroC, Inc.

guaranteed that all entries in the evictor will be idle and therefore will be evicted.

Java

```
synchronized public final void
deactivate(String category)
{
    _size = 0;
    evictServants();
}
```

Note that, with this implementation of `evictServants`, we only scan the tail section of the evictor queue for servants to evict. If we have long-running operations, this allows the number of servants in the queue to remain above the evictor size if the servants in the tail section have a non-zero use count. This means that, even immediately after calling `evictServants`, the queue length can still exceed the evictor size.

We can adopt a more aggressive strategy for eviction: instead of scanning only the excess entries in the queue, if, after looking in the tail section of the queue, we still have more servants in the queue than the queue size, we keep scanning for servants with a zero use count until the queue size drops below the limit. This alternative version of `evictServants` looks as follows:

Java

```
private void evictServants()
{
    //
    // If the evictor queue has grown larger than the limit,
    // look at the excess elements to see whether any of them
    // can be evicted.
    //
    java.util.Iterator<Ice.Identity> p = _queue.riterator();
    int numEntries = _map.size();
    for (int i = 0; i < excessEntries && _map.size() > _size;
         ++i) {
        Ice.Identity id = p.next();
        EvictorEntry e = _map.get(id);
        if (e.useCount == 0) {
            evict(e.servant, e.userCookie); // Down-call
            e.queuePos.remove();
            _map.remove(id);
        }
    }
}
```

The only difference in this version is that the terminating condition for the `for`-loop has changed: instead of scanning only the excess entries for servants with a use count, this version keeps scanning until the evictor size drops below the limit.

Which version is more appropriate depends on your application: if locating and evicting servants is expensive, and memory is not at a premium, the first version (which only scans the tail section) is more appropriate; if you want to keep memory consumption to a minimum, the second version is more appropriate. Also keep in mind that the difference between the two versions is significant only if you have long-running operations and many concurrent invocations from clients; otherwise, there is no point in more aggressively scanning for servants to remove because they are going be become idle again very quickly and get evicted as soon as the next request arrives.

Using Servant Evictors in Java

Copyright 2016, ZeroC, Inc.

Using a servant evictor is simply a matter of deriving a class from `EvictorBase` and implementing the `add` and `evict` methods. You can turn a servant locator into an evictor by simply taking the code that you wrote for `locate` and placing it into `add` — `EvictorBase` then takes care of maintaining the cache in least-recently used order and evicting servants as necessary. Unless you have clean-up requirements for your servants (such as closing network connections or database handles), the implementation of `evict` can be left empty.

One of the nice aspects of evictors is that you do not need to change anything in your servant implementation: the servants are ignorant of the fact that an evictor is in use. This makes it very easy to add an evictor to an already existing code base with little disturbance of the source code.

Evictors can provide substantial performance improvements over default servants: especially if initialization of servants is expensive (for example, because servant state must be initialized by reading from a network), an evictor performs much better than a default servant, while keeping memory requirements low.

See Also

- Servant Evictors
- Object Identity
- Default Servants

Copyright 2016, ZeroC, Inc.

Implementing a Servant Evictor in C-Sharp

On this page:

- A Linked List to Support Eviction in C#
- The EvictorBase Class in C#
- Using Servant Evictors in C#

A Linked List to Support Eviction in C#

The `System.Collections` classes do not provide a container that does not invalidate iterators when we modify the contents of the container but, to efficiently implement an evictor, we need such a container. To deal with this, we use a special-purpose linked list implementation, `Evictor.LinkedList`, that does not invalidate iterators when we delete or add an element. For brevity, we only show the interface of `LinkedList` here — you can find the implementation in the code examples for this manual in the Ice distribution.

C#

```
namespace Evictor
{
    public class LinkedList<T> : ICollection<T>, ICollection, ICloneable
    {
        public LinkedList();

        public int Count { get; }

        public void Add(T value);
        public void AddFirst(T value);
        public void Clear();
        public bool Contains(T value);
        public bool Remove(T value);

        public IEnumerator GetEnumerator();

        public class Enumerator : IEnumerator<T>, IEnumerator, IDisposable
        {
            public void Reset();

            public T Current { get; }

            public bool MoveNext();
            public bool MovePrev();
            public void Remove();
            public void Dispose();
        }

        public void CopyTo(T[] array, int index);
        public void CopyTo(Array array, int index);

        public object Clone();

        public bool IsReadOnly { get; }
        public bool IsSynchronized { get; }
        public object SyncRoot { get; }
    }
}
```


Copyright 2016, ZeroC, Inc.

The Add method appends an element to the list, and the AddFirst method prepends an element to the list. GetEnumerator returns an enumerator for the list elements; immediately after calling GetEnumerator, the enumerator does not point at any element until you call either MoveNext or MovePrev, which position the enumerator at the first and last element, respectively. Current returns the element at the enumerator position, and Remove deletes the element at the current position and leaves the enumerator pointing at no element. Calling MoveNext or MovePrev after calling Remove positions the enumerator at the element following or preceding the deleted element, respectively. MoveNext and MovePrev return true if they have positioned the enumerator on an element; otherwise, they return false and leave the enumerator position on the last and first element, respectively.

The EvictorBase Class in C#

Given this LinkedList, we can implement the evictor. The evictor we show here is designed as an abstract base class: in order to use it, you derive a class from the Evictor.EvictorBase base class and implement two methods that are called by the evictor when it needs to add or evict a servant. This leads to a class definition as follows:

Copyright 2016, ZeroC, Inc.

C#

```
namespace Evictor
{
    public abstract class EvictorBase : Ice.ServantLocator
    {
        public EvictorBase()
        {
            _size = 1000;
        }

        public EvictorBase(int size)
        {
            _size = size < 0 ? 1000 : size;
        }

        protected abstract Ice.Object add(Ice.Current c, out object cookie);

        protected abstract void evict(Ice.Object servant, object cookie);

        public Ice.Object locate(Ice.Current c, out object cookie)
        {
            lock(this)
            {
                // ...
            }
        }

        public void finished(Ice.Current c, Ice.Object o, object cookie)
        {
            lock(this)
            {
                // ...
            }
        }

        public void deactivate(string category)
        {
            lock(this)
            {
                // ...
            }
        }

        private int _size;
    }
}
```

Note that the evictor has constructors to set the size of the queue, with a default size of 1000.

The `locate`, `finished`, and `deactivate` methods are inherited from the `ServantLocator` base class; these methods implement the logic to maintain the queue in LRU order and to add and evict servants as needed. The methods use a `lock(this)` statement for their body, so the evictor's internal data structures are protected from concurrent access.

Copyright 2016, ZeroC, Inc.

The `add` and `evict` methods are called by the evictor when it needs to add a new servant to the queue and when it evicts a servant from the queue. Note that these functions are abstract, so they must be implemented in a derived class. The job of `add` is to instantiate and initialize a servant for use by the evictor. The `evict` function is called by the evictor when it evicts a servant, allowing the subclass to perform any cleanup. Note that `add` can return a cookie that the evictor passes to `evict`, so you can move context information from `add` to `evict`.

Next, we need to consider the data structures that are needed to support our evictor implementation. We require two main data structures:

1. A map that maps object identities to servants, so we can efficiently decide whether we have a servant for an incoming request in memory or not.
2. A list that implements the evictor queue. The list is kept in LRU order at all times.

The evictor map not only stores servants but also keeps track of some administrative information:

1. The map stores the cookie that is returned from `add`, so we can pass that same cookie to `evict`.
2. The map stores an iterator into the evictor queue that marks the position of the servant in the queue.
3. The map stores a use count that is incremented whenever an operation is dispatched into a servant, and decremented whenever an operation completes.

The last two points deserve some extra explanation.

- The evictor queue must be maintained in least-recently used order, that is, every time an invocation arrives and we find an entry for the identity in the evictor map, we also must locate the servant's identity on the evictor queue and move it to the front of the queue. However, scanning for that entry is inefficient because it requires *O*(*n*) time. To get around this, we store an iterator in the evictor map that marks the corresponding entry's position in the evictor queue. This allows us to dequeue the entry from its current position and enqueue it at the head of the queue in *O*(1) time, using the `Evictor.LinkedList` implementation.

- We maintain a use count as part of the map in order to avoid incorrect eviction of servants. Suppose a client invokes a long-running operation on an Ice object with identity *I*. In response, the evictor adds a servant for *I* to the evictor queue. While the original invocation is still executing, other clients invoke operations on various Ice objects, which leads to more servants for other object identities being added to the queue. As a result, the servant for identity *I* gradually migrates toward the tail of the queue. If enough client requests for other Ice objects arrive while the operation on object *I* is still executing, the servant for *I* could be evicted while it is still executing the original request.

  By itself, this will not do any harm. However, if the servant is evicted and a client then invokes another request on object *I*, the evictor would have no idea that a servant for *I* is still around and would add a second servant for *I*. However, having two servants for the same Ice object in memory is likely to cause problems, especially if the servant's operation implementations write to a database.

  The use count allows us to avoid this problem: we keep track of how many requests are currently executing inside each servant and, while a servant is busy, avoid evicting that servant. As a result, the queue size is not a hard upper limit: long-running operations can temporarily cause more servants than the limit to appear in the queue. However, as soon as excess servants become idle, they are evicted as usual.

Finally, our `locate` and `finished` implementations will need to exchange a cookie that contains a reference to the entry in the evictor map. This is necessary so that `finished` can decrement the servant's use count.

This leads to the following definitions in the private section of our evictor:

Copyright 2016, ZeroC, Inc.

C#

```
namespace Evictor
{
    using System.Collections.Generic;

    public abstract class EvictorBase : Ice.ServantLocator
    {
        // ...

        private class EvictorEntry
        {
            internal Ice.Object servant;
            internal object userCookie;
            internal LinkedList<Ice.Identity>.Enumerator queuePos;
            internal int useCount;
        }

        private void evictServants()
        {
            // ...
        }

        private Dictionary<Ice.Identity, EvictorEntry> _map =
            new Dictionary<Ice.Identity, EvictorEntry>();
        private LinkedList<Ice.Identity> _queue =
            new LinkedList<Ice.Identity>();
        private int _size;
    }
}
```

Note that the evictor stores the evictor map, queue, and the queue size in the private data members `_map`, `_queue`, and `_size`. The map key is the identity of the Ice object, and the lookup value is of type `EvictorEntry`. The queue simply stores identities, of type `Ice.Identity`.

The `evictServants` member function takes care of evicting servants when the queue length exceeds its limit — we will discuss this function in more detail shortly.

Almost all the action of the evictor takes place in the implementation of `locate`:

Copyright 2016, ZeroC, Inc.

C#

```
public Ice.Object locate(Ice.Current c, out object cookie)
{
    lock(this)
    {
        //
        // Check if we a servant in the map already.
        //
        EvictorEntry entry = _map[c.id];
        if (entry != null) {
            //
            // Got an entry already, dequeue the entry from
            // its current position.
            //
            entry.queuePos.Remove();
        } else {
            //
            // We do not have an entry. Ask the derived class to
            // instantiate a servant and add an entry to the map.
            //
            entry = new EvictorEntry();
            entry.servant = add(c, out entry.userCookie);
            if (entry.servant == null) {
                cookie = null;
                return null;
            }
            entry.useCount = 0;
            _map[c.id] = entry;
        }

        //
        // Increment the use count of the servant and enqueue
        // the entry at the front, so we get LRU order.
        //
        ++(entry.useCount);
        _queue.AddFirst(c.id);
        entry.queuePos = (LinkedList<Ice.Identity>.Enumerator)_queue.GetEnumerator(
);
        entry.queuePos.MoveNext();

        cookie = entry;

        return entry.servant;
    }
}
```

The code uses an `EvictorEntry` as the cookie that is returned from `locate` and will be passed by the Ice run time to the corresponding call to `finished`.

We first look for an existing entry in the evictor map, using the object identity as the key. If we have an entry in the map already, we dequeue the corresponding identity from the evictor queue. (The `queuePos` member of `EvictorEntry` is an iterator that marks that entry's position in the evictor queue.)

Copyright 2016, ZeroC, Inc.

Otherwise, we do not have an entry in the map, so we create a new one and call the `add` method. This is a down-call to the concrete class that will be derived from `EvictorBase`. The implementation of `add` must attempt to locate the object state for the Ice object with the identity passed inside the `Current` object and either return a servant as usual, or return null or throw an exception to indicate failure. If `add` returns null, we return null to let the Ice run time know that no servant could be found for the current request. If `add` succeeds, we initialize the entry's use count to zero and insert the entry into the evictor map.

The final few lines of code increment the entry's use count, add the entry at the head of the evictor queue, store the entry's position in the queue, and initialize the cookie that is returned from `locate`, before returning the servant to the Ice run time.

The implementation of `finished` is comparatively simple. It decrements the use count of the entry and then calls `evictServants` to get rid of any servants that might need to be evicted:

C#

```
public void finished(Ice.Current c, Ice.Object o, object cookie)
{
    lock(this)
    {
        EvictorEntry entry = (EvictorEntry)cookie;

        //
        // Decrement use count and check if
        // there is something to evict.
        //
        --(entry.useCount);
        evictServants();
    }
}
```

In turn, `evictServants` examines the evictor queue: if the queue length exceeds the evictor's size, the excess entries are scanned. Any entries with a zero use count are then evicted:

Copyright 2016, ZeroC, Inc.

C#

```
private void evictServants()
{
    //
    // If the evictor queue has grown larger than the limit,
    // look at the excess elements to see whether any of them
    // can be evicted.
    //
    LinkedList<Ice.Identity>.Enumerator p =
        (LinkedList<Ice.Identity>.Enumerator)_queue.GetEnumerator();
    int excessEntries = _map.Count - _size;
    for (int i = 0; i < excessEntries; ++i) {
        p.MovePrev();
        Ice.Identity id = p.Current;
        EvictorEntry e = _map[id];
        if (e.useCount == 0) {
            evict(e.servant, e.userCookie); // Down-call
            p.Remove();
            _map.Remove(id);
        }
    }
}
```

The code scans the excess entries, starting at the tail of the evictor queue. If an entry has a zero use count, it is evicted: after calling the `evict` member function in the derived class, the code removes the evicted entry from both the map and the queue.

Finally, the implementation of `deactivate` sets the evictor size to zero and then calls `evictServants`. This results in eviction of all servants. The Ice run time guarantees to call `deactivate` only once no more requests are executing in an object adapter; as a result, it is guaranteed that all entries in the evictor will be idle and therefore will be evicted.

C#

```
public void deactivate(string category)
{
    lock(this)
    {
        _size = 0;
        evictServants();
    }
}
```

Note that, with this implementation of `evictServants`, we only scan the tail section of the evictor queue for servants to evict. If we have long-running operations, this allows the number of servants in the queue to remain above the evictor size if the servants in the tail section have a non-zero use count. This means that, even immediately after calling `evictServants`, the queue length can still exceed the evictor size.

We can adopt a more aggressive strategy for eviction: instead of scanning only the excess entries in the queue, if, after looking in the tail section of the queue, we still have more servants in the queue than the queue size, we keep scanning for servants with a zero use count until the queue size drops below the limit. This alternative version of `evictServants` looks as follows:

Copyright 2016, ZeroC, Inc.

C#

```
private void evictServants()
{
    //
    // If the evictor queue has grown larger than the limit,
    // look at the excess elements to see whether any of them
    // can be evicted.
    //
    LinkedList<Ice.Identity>.Enumerator p =
        (LinkedList<Ice.Identity>.Enumerator)_queue.GetEnumerator();
    int numEntries = _map.Count;
    for (int i = 0; i < numEntries && _map.Count > _size; ++i) {
        p.MovePrev();
        Ice.Identity id = p.Current;
        EvictorEntry e = _map[id];
        if (e.useCount == 0) {
            evict(e.servant, e.userCookie); // Down-call
            p.Remove();
            _map.Remove(id);
        }
    }
}
```

The only difference in this version is that the terminating condition for the `for`-loop has changed: instead of scanning only the excess entries for servants with a use count, this version keeps scanning until the evictor size drops below the limit.

Which version is more appropriate depends on your application: if locating and evicting servants is expensive, and memory is not at a premium, the first version (which only scans the tail section) is more appropriate; if you want to keep memory consumption to a minimum, the second version is more appropriate. Also keep in mind that the difference between the two versions is significant only if you have long-running operations and many concurrent invocations from clients; otherwise, there is no point in more aggressively scanning for servants to remove because they are going be become idle again very quickly and get evicted as soon as the next request arrives.

Using Servant Evictors in C#

Using a servant evictor is simply a matter of deriving a class from `EvictorBase` and implementing the `add` and `evict` methods. You can turn a servant locator into an evictor by simply taking the code that you wrote for `locate` and placing it into `add` — `EvictorBase` then takes care of maintaining the cache in least-recently used order and evicting servants as necessary. Unless you have clean-up requirements for your servants (such as closing network connections or database handles), the implementation of `evict` can be left empty.

One of the nice aspects of evictors is that you do not need to change anything in your servant implementation: the servants are ignorant of the fact that an evictor is in use. This makes it very easy to add an evictor to an already existing code base with little disturbance of the source code.

Evictors can provide substantial performance improvements over default servants: especially if initialization of servants is expensive (for example, because servant state must be initialized by reading from a network), an evictor performs much better than a default servant, while keeping memory requirements low.

See Also

- Servant Evictors
- Object Identity
- Default Servants

Copyright 2016, ZeroC, Inc.

Object Life Cycle

Topics

- Understanding Object Life Cycle
- Object Existence and Non-Existence
- Life Cycle of Proxies, Servants, and Ice Objects
- Object Creation
- Object Destruction
- Removing Cyclic Dependencies
- Object Identity and Uniqueness
- Object Life Cycle for the File System Application
- Avoiding Server-Side Garbage

Copyright 2016, ZeroC, Inc.

Understanding Object Life Cycle

Object life cycle generally refers to how an object-oriented application (whether distributed or not) creates and destroys objects. For distributed applications, life cycle management presents particular challenges. For example, destruction of objects often can be surprisingly complex, especially in threaded applications. Before we go into the details of object creation and destruction, we need to have a closer look what we mean by the terms "life cycle" and "object" in this context.

Object life cycle refers to the act of creation and destruction of objects. For example, with our file system application, we may start out with an empty file system that only contains a root directory. Over time, clients (by as yet unspecified means) add new directories and files to the file system. For example, a client might create a new directory called `MyPoems` underneath the root directory. Some time later, the same or a different client might decide to remove this directory again, returning the file system to its previous empty state. This pattern of creation and destruction is known as object life cycle.

The life cycle of distributed objects raises a number of interesting and challenging questions. For example, what should happen if a client destroys a file while another client is reading or writing that file? And how do we prevent two files with the same name from existing in the same directory? Another interesting scenario is illustrated by the following sequence of events:

1. Client A creates a file called `DraftPoem` in the root directory and uses it for a while.
2. Some time later, client B destroys the `DraftPoem` file so it no longer exists.
3. Some time later still, client C creates a new `DraftPoem` file in the root directory, with different contents.
4. Finally, client A attempts to access the `DraftPoem` file it created earlier.

What should happen when, in the final step, client A tries to use the `DraftPoem` file? Should the client's attempt succeed and simply operate on the new contents of the file that were placed there by client C? Or should client A's attempt fail because, after all, the new `DraftPoem` file is, in a sense, a completely different file from the original one, even though it has the same name?

The answers to such questions cannot be made in general. Instead, meaningful answers depend on the semantics that each individual application attaches to object life cycle. In this chapter, we will explore the various possible interpretations and how to implement them correctly, particularly for threaded applications.

See Also

- Slice for a Simple File System

Copyright 2016, ZeroC, Inc.

Object Existence and Non-Existence

Before we talk about how to create and destroy objects, we need to look at a more basic concept, namely that of object existence. What does it mean for an object to "exist" and, more fundamentally, what do we mean by the term "object"?

As mentioned in Ice Objects, an *Ice object* is a conceptual entity, or abstraction that does not really exist. On the client side, the concrete representation of an Ice object is a proxy and, on the server side, the concrete representation of an Ice object is a servant. Proxies and servants are the concrete programming-language artifacts that represent Ice objects.

Because Ice objects are abstract, conceptual entities, they are invisible to the Ice run time and to the application code — only proxies and servants are real and visible. It follows that, to determine whether an *Ice* object exists, any determination must rely on proxies and servants, because they are the only tangible entities in the system.

On this page:

- Object Non-Existence
- Object Existence
- Indeterminate Object State
- Authoritative Object Existence Semantics

Object Non-Existence

Here is the definitive statement of what it means for an Ice object to *not* exist:

> *An Ice object does not exist if an invocation on the object raises an `ObjectNotExistException`.*

This may seem self-evident but, on closer examination, is a little more subtle than you might expect. In particular, Ice object existence has meaning *only within the context of a particular invocation*. If that invocation raises `ObjectNotExistException`, the object is known to not exist. Note that this says nothing about whether concurrent or future requests to that object will also raise `ObjectNotExistException` — they may or may not, depending on the semantics that are implemented by the application.

Also note that, because all the Ice run time knows about are servants, an `ObjectNotExistException` really indicates that a *servant* for the request could not be found at the time the request was made. This means that, ultimately, it is the application that attaches the meaning "the Ice object does not exist" to this exception.

In theory, the application can attach any meaning it likes to `ObjectNotExistException` and a server can throw this exception for whatever reason it sees fit; in practice, however, we recommend that you do not do this because it breaks with existing convention and is potentially confusing. You should reserve this exception for its intended meaning and not abuse it for other purposes.

Object Existence

The preceding definition does not say anything about object existence if something other than `ObjectNotExistException` is returned in response to a particular request. So, here is the definitive statement of what it means for an Ice object to exist:

> *An Ice object exists if a twoway invocation on the object either succeeds, raises a user exception, or raises `FacetNotExistException` or `OperationNotExistException`.*

It is self-evident that an Ice object exists if a twoway invocation on it succeeds: obviously, the object received the invocation, processed it, and returned a result. However, note the qualification: this is true only for *twoway* invocations; for oneway and datagram invocations, nothing can be inferred about the existence of the corresponding Ice object by invoking an operation on it: because there is no reply from the server, the client-side Ice run time has no idea whether the request was dispatched successfully in the server or not. This includes user exceptions, `ObjectNotExistException`, `FacetNotExistException`, and `OperationNotExistException` — these exceptions are never raised by oneway and datagram invocations, regardless of the actual state of the target object.

If a twoway invocation raises a user exception, the Ice object obviously exists: the Ice run time never raises user exceptions so, for an invocation to raise a user exception, the invocation was dispatched successfully in the server, and the operation implementation in the servant raised the exception.

If a twoway invocation raises `FacetNotExistException`, we do know that the corresponding Ice object indeed exists: the Ice run time raises `FacetNotExistException` only if it can find the identity of the target object in the Active Servant Map (ASM), but cannot find the facet that was specified by the client.

> Note that, if you use servant locators, for these semantics to hold, your servant locator must correctly raise `FacetNotExistException` in the `locate` operation (instead of returning null or raising `ObjectNotExistException`) if an Ice object exists, but the particular target facet does not exist.

As a corollary to the preceding two definitions, we can state:

Copyright 2016, ZeroC, Inc.

> *A facet does not exist if a twoway invocation on the object raises `ObjectNotExistException` or `FacetNotExistException`.*
>
> *A facet exists if a twoway invocation on the object either succeeds, or raises `OperationNotExistException`.*

These definitions simply capture the fact that a facet is a "sub-object" of an Ice object: if an invocation raises `ObjectNotExistException`, we know that the facet does not exist either because, for a facet to exist, its Ice object must exist.

If an operation raises `OperationNotExistException`, we know that both the target Ice object and the target facet exist. However, the operation that the client attempted to invoke does not. (This is possible only if you use dynamic invocation or if you have mis-matched Slice definitions for client and server.)

Indeterminate Object State

The preceding definitions clearly state under what circumstances we can conclude that an Ice object (or its facet) does or does not exist. However, the preceding definitions are incomplete because operation invocations can have outcomes other than success or failure with `ObjectNotExistException`, `FacetNotExistException`, or `OperationNotExistException`. For example, a client might receive a `MarshalException`, `UnknownLocalException`, `UnknownException`, or `TimeoutException`. In that case, the client cannot draw any conclusions about whether the Ice object on which it invoked a twoway operation exists or not — the exceptions simply indicate that something went wrong while the invocation was processed. So, to complete our definitions, we can state:

> *If a twoway invocation raises a run-time exception other than `ObjectNotExistException`, `FacetNotExistException`, or `OperationNotExistException`, nothing is known about the existence or non-existence of the Ice object that was the target of the invocation. Furthermore, it is impossible to determine the state of existence of an Ice object with a oneway or datagram invocation.*

Authoritative Object Existence Semantics

The preceding definitions capture the fact that, to make a determination of object existence or non-existence, the client-side Ice run time must be able to contact the server and, moreover, receive a reply from the server:

- If the server can be contacted and returns a successful reply for an invocation, the Ice object exists.
- If the server can be contacted and returns an `ObjectNotExistException` (or `FacetNotExistException`), the Ice object (or facet) does not exist. If the server returns an `OperationNotExistException`, the Ice object (and its facet) exists, but does not provide the requested operation, which indicates a type mismatch due to client and server using out-of-sync Slice definitions or due to incorrect use of dynamic invocation.
- If the server cannot be contacted, does not return a reply (as for oneway and datagram invocations), or if anything at all goes wrong with the process of sending an invocation, processing it in the server, and returning the reply, nothing is known about the state of the Ice object, including its existence or non-existence.

Another way of looking at this is that a decision as to whether an object exists or not is *never* made by the Ice run time and, instead, is *always* made by the server-side *application* code:

- If an invocation completes successfully, the server-side application code was involved because it processed the invocation.
- If an invocation returns `ObjectNotExistException` or `FacetNotExistException`, the server-side application code was also involved:
  - either the Ice run time could not find a servant for the invocation in the ASM, in which case the application code was involved by virtue of not having added a servant to the ASM in the first place, or
  - the Ice run time consulted a servant locator that explicitly returned null or raised `ObjectNotExistException` or `FacetNotExistException`.

This means that `ObjectNotExistException` and `FacetNotExistException` are *authoritative*: when you receive these exceptions, you can always believe what they tell you — the Ice run time never raises these exceptions without consulting your code, either implicitly (via an ASM lookup) or explicitly (by calling a servant locator's `locate` operation).

These semantics are motivated by the need to keep the Ice run time stateless with respect to object existence. For example, it would be nice to have stronger semantics, such as a promise that "once an Ice object has existed and been destroyed, all future requests to that Ice object also raise `ObjectNotExistException`". However, to implement these semantics, the Ice run time would have to remember all object identities that were used in the past, and prevent their reuse for new Ice objects. Of course, this would be inherently non-scalable. In addition, it would prevent applications from controlling object identity; allowing such control for applications is important however, for example, to link the identity of an Ice object to its persistent state in a database.

Note that, if the implementation of an operation calls another operation, dealing with `ObjectNotExistException` may require some care. For example, suppose that the client holds a proxy to an object of type `Service` and invokes an operation `provideService` on it:

Copyright 2016, ZeroC, Inc.

C++

```
ServicePrx service = ...;

try {
    service->provideService();
} catch (const ObjectNotExistException&) {
    // Service does not exist.
}
```

Here is the implementation of `provideService` in the server, which makes a call on a helper object to implement the operation:

C++

```
void
ServiceI::provideService(const Ice::Current&)
{
    // ...
    proxyToHelper->someOp();
    // ...
}
```

If `proxyToHelper` happens to point at an object that was destroyed previously, the call to `someOp` will throw `ObjectNotExistException`. If the implementation of `provideService` does not intercept this exception, the exception will propagate all the way back to the client, who will conclude that the service has been destroyed when, in fact, the service still exists but the helper object used to implement `provideService` no longer exists.

Usually, this scenario is not a serious problem. Most often, the helper object cannot be destroyed while it is needed by `provideService` due to the way the application is structured. In that case, no special action is necessary because `someOp` will never throw `ObjectNotExistException`. On the other hand, if it is possible for the helper object to be destroyed, `provideService` can wrap a try-catch block for `ObjectNotExistException` around the call to `someOp` and throw an appropriate user exception from the exception handler (such as `ResourceUnavailable` or similar).

See Also

- Terminology
- Servant Locators
- Oneway Invocations
- Datagram Invocations
- Versioning
- Dynamic Ice
- Freeze

Copyright 2016, ZeroC, Inc.

Life Cycle of Proxies, Servants, and Ice Objects

It is important to be aware of the different roles of proxies, servants, and Ice objects in a system. Proxies are the client-side representation of Ice objects and servants are the server-side representation of Ice objects. Proxies, servants, and Ice objects have completely independent life cycles. Clients can create and destroy proxies with or without a corresponding servant or Ice object in existence, servers can create and destroy servants with or without a corresponding proxy or Ice object in existence and, most importantly, Ice objects can exist or not exist regardless of whether corresponding proxies or servants exist. Here are a few examples to illustrate this:

C++

```
{
    Ice::ObjectPtr obj = communicator->stringToProxy("hello:tcp -p 10000");
    // Proxy exists now.

} // Proxy ceases to exist.
```

This code creates a proxy to an Ice object with the identity `Hello`. The server for this Ice object is expected to listen for invocations on the same host as the client, on port 10000, using the TCP/IP protocol. The proxy exists as soon as the call to `stringToProxy` completes and, thereafter, can by used by the client to make invocations on the corresponding Ice object.

However, note that this code says nothing at all about whether or not the corresponding Ice object exists. In particular, there might not be any Ice object with the identity `Hello`. Or there might be such an object, but the server for it may be down or unreachable. It is only when the client makes an invocation on the proxy that we get to find out whether the object exists, does not exist, or cannot be reached.

Similarly, at the end of the scope enclosing the `obj` variable in the preceding code, the proxy goes out of scope and is destroyed. Again, this says nothing about the state of the corresponding Ice object or its servant. This shows that the life cycle of a proxy is completely independent of the life cycle of its Ice object and the servant for that Ice object: clients can create and destroy proxies whenever they feel like it, and doing so has no implications for Ice objects or servant creation or destruction.

Here is another code example, this time for the server side:

C++

```
{
    FileIPtr file = new FileI("DraftPoem", root);
    // Servant exists now.

} // Servant ceases to exist.
```

Here, the server instantiates a servant for a File object by creating a `FileI` instance. The servant comes into being as soon as the call to `new` completes and ceases to exist as soon as the scope enclosing the `file` variable closes. Note that, as for proxies, the life cycle of the servant is completely independent of the life cycle of proxies and Ice objects. Clearly, the server can create and destroy a servant regardless of whether there are any proxies in existence for the corresponding Ice object.

Similarly, an Ice object can exist even if no servants exist for it. For example, our Ice objects might be persistent and stored in a database; in that case, if we switch off the server for our Ice objects, no servants exist for these Ice objects, even though the Ice objects continue to exist — the Ice objects are temporarily inaccessible, but exist regardless and, once their server is restarted, will become accessible again.

Conversely, a servant can exist without its corresponding Ice object. The mere creation of a servant does nothing, as far as the Ice run time is concerned. It is only once a servant is added to the Active Servant Map (ASM) (or a servant locator returns the servant from `locate`) that the servant incarnates its Ice object.

Finally, an Ice object can exist independently of proxies and servants. For example, returning to the database example, we might have an Ice server that acts as a front end to an online telephone book: each entry in the phone book corresponds to a separate Ice object. When a client invokes an operation, the server uses the identity of the incoming request to determine which Ice object is the target of the request and

Copyright 2016, ZeroC, Inc.

then contacts the back-end database to, for example, return the street address of the entry. With such a design, entries can be added to and removed from the back-end database quite independently of what happens to proxies and servants — the server finds out whether an Ice object exists only when it accesses the back-end database.

The only time that the life cycle of an Ice object and a servant are linked is during an invocation on that Ice object: for an invocation to complete successfully, a servant must exist for the *duration of the invocation*. What happens to the servant thereafter is irrelevant to clients and, in general, is irrelevant to the corresponding Ice object.

It is important to be clear about the independence of the life cycles of proxies, servants, and Ice objects because this independence has profound implications for how you need to implement object life cycle. In particular, to destroy an Ice object, a client cannot simply destroy its proxy for an object because the server is completely unaware when a client does this.

Distributed object systems such as DCOM implement these semantics. However, this design is inherently non-scalable because of the cost of globally tracking proxy creation and destruction.

See Also

- The Active Servant Map
- Servant Locators
- Object Identity

Copyright 2016, ZeroC, Inc.

Object Creation

Now that we understand what it means for an Ice object to exist, we can look at what is involved in creating an Ice object. Fundamentally, there is only one way for an Ice object to come into being: the server must instantiate a servant for the object and add an entry for that servant to the Active Servant Map (ASM) (or, alternatively, arrange for a servant locator to return a servant from its `locate` operation).

For the remainder of this chapter, we will ignore the distinction between using the ASM and a servant locator and simply assume that the code uses the ASM. This is because servant locators do not alter the discussion: if `locate` returns a servant, that is the same as a successful lookup in the ASM; if `locate` returns null or throws `ObjectNotExistException`, that is the same as an unsuccessful lookup in the ASM.

One obvious way for a server to create a servant is to, well, simply instantiate it and add it to the ASM of its own accord. For example:

C++

```
DirectoryIPtr root = new DirectoryI("/", 0);
adapter->addWithUUID(root); // Ice object exists now
```

The servant exists as soon as the call to `new` completes, and the Ice object exists as soon as the code adds the servant to the ASM: at that point, the Ice object becomes reachable to clients who hold a proxy to it.

This is the way we created Ice objects for our file system application earlier in the manual. However, doing so is not all that interesting because the only files and directories that exist are those that the server decides to create when it starts up. What we really want is a way for *clients* to create and destroy directories and files.

On this page:

- Creating an Object with a Factory
- Implementing a Factory Operation

Creating an Object with a Factory

The canonical way to create an object is to use the factory pattern [1]. The factory pattern, in a nutshell, says that objects are created by invoking an operation (usually called `create`) on an object factory:

Slice

```
interface PhoneEntry {
    idempotent string name();
    idempotent string getNumber();
    idempotent void setNumber(string phNum);
};

exception PhoneEntryExists {
    string name;
    string phNum;
};

interface PhoneEntryFactory {
    PhoneEntry* create(string name, string phNum)
        throws PhoneEntryExists;
};
```


Copyright 2016, ZeroC, Inc.

Rather than continue with the file system example, we will simplify the discussion for the time being by using the phone book example mentioned earlier; we will return to the file system application to explore more complex issues.

The entries in the phone book consist of simple name-number pairs. The interface to each entry is called `PhoneEntry` and provides operations to read the name and to read and write the phone number. (For a real application, the objects would likely be more complex and encapsulate more state. However, these simple objects will do for the purposes of this discussion.)

To create a new entry, a client calls the `create` operation on a `PhoneEntryFactory` object. (The factory is a singleton object [1] — that is, only one instance of that interface exists in the server.) It is the job of `create` to create a new `PhoneEntry` object, using the supplied name as the object identity.

An immediate consequence of using the name as the object identity is that `create` can raise a `PhoneEntryExists` exception: presumably, if a client attempts to create an entry with the same name as an already-existing entry, we need to let the client know about this. (Whether this is an appropriate design is something we examine more closely in Object Identity and Uniqueness.)

`create` returns a proxy to the newly-created object, so the client can use that proxy to invoke operations. However, this is by convention only. For example, `create` could be a `void` operation if the client has some other way to eventually get a proxy to the new object (such as creating the proxy from a string, or locating the proxy via a search operation). Alternatively, you could define "bulk" creation operations that allow clients to create several new objects with a single RPC. As far as the Ice run time is concerned, there is nothing special about a factory operation: a factory operation is just like any other operation; it just so happens that a factory operation creates a new Ice object as a side effect of being called, that is, the *implementation* of the operation is what creates the object, not the Ice run time.

Also note that `create` accepts a `name` and a `phNum` parameter, so it can initialize the new object. This is not compulsory, but generally a good idea. An alternate factory operation could be:

Slice

```
interface PhoneEntryFactory {
    PhoneEntry* create(string name)
        throws PhoneEntryExists;
};
```

With this design, the assumption is that the client will call `setNumber` after it has created the object. However, in general, allowing objects that are not fully initialized is a bad idea: it all too easily happens that a client either forgets to complete the initialization, or happens to crash or get disconnected before it can complete the initialization. Either way, we end up with a partially-initialized object in the system that can cause surprises later.

This is the approach taken by COM's `CoCreateObject`, which suffers from just that problem.

Similarly, so-called `generic` factories are also something to be avoided:

Slice

```
dictionary<string, string> Params;

exception CannotCreateException {
    string reason;
};

interface GenericFactory {
    Object* create(Params p)
        throws CannotCreateException;
};
```


Copyright 2016, ZeroC, Inc.

The intent here is that a `GenericFactory` can be used to create any kind of object; the `Params` dictionary allows an arbitrary number of parameters to be passed to the `create` operation in the form of name — value pairs, for example:

C++

```
GenericFactoryPrx factory = ...;

Ice::ObjectPrx obj;
Params p;

// Make a car.
//
p["Make"] = "Ford";
p["Model"] = "Falcon";
obj = factory->create(p);
CarPrx car = CarPrx::checkedCast(obj);

// Make a horse.
//
p.clear();
p["Breed"] = "Clydesdale";
p["Sex"] = "Male";
obj = factory->create(p);
HorsePrx horse = HorsePrx::checkedCast(obj);
```

We strongly discourage you from creating factory interfaces such as this, unless you have a good overriding reason: generic factories undermine type safety and are much more error-prone than strongly-typed factories.

Implementing a Factory Operation

The implementation of an object factory is simplicity itself. Here is how we could implement the `create` operation for our `PhoneEntryFactory`:

C++

```
PhoneEntryPrx
PhoneEntryFactory::create(const string& name, const string& phNum,
const Current& c)
{
    try {
        CommunicatorPtr comm = c.adapter->getCommunicator();
        PhoneEntryPtr servant = new PhoneEntryI(name, phNum);
        return PhoneEntryPrx::uncheckedCast(c.adapter->add(servant,
comm->stringToIdentity(name)));
    } catch (const Ice::AlreadyRegisteredException&) {
        throw PhoneEntryExists(name, phNum);
    }
}
```

The `create` function instantiates a new `PhoneEntryI` object (which is the servant for the new `PhoneEntry` object), adds the servant to the ASM, and returns the proxy for the new object. Adding the servant to the ASM is what creates the new Ice object, and client requests are

Copyright 2016, ZeroC, Inc.

dispatched to the new object as soon as that entry appears in the ASM (assuming the object adapter is active).

Note that, even though this code contains no explicit lock, it is thread-safe. The `add` operation on the object adapter is atomic: if two clients concurrently add a servant with the same identity, exactly one thread succeeds in adding the entry to the ASM; the other thread receives an `AlreadyRegisteredException`. Similarly, if two clients concurrently call `create` for different entries, the two calls execute concurrently in the server (if the server is multi-threaded); the implementation of `add` in the Ice run time uses appropriate locks to ensure that concurrent updates to the ASM cannot corrupt anything.

See Also

- The Active Servant Map
- Servant Locators
- Object Identity
- Object Identity and Uniqueness
- Object Life Cycle for the File System Application
- Object Adapter States
- Servant Activation and Deactivation

References

1. Gamma, E., et al. 1994. Design Patterns. Reading, MA: Addison-Wesley.

Copyright 2016, ZeroC, Inc.

Object Destruction

Now that clients can create `PhoneEntry` objects, let us consider how to allow clients to destroy them again. One obvious design is to add a `destroy` operation to the factory — after all, seeing that a factory knows how to create objects, it stands to reason that it also knows how to destroy them again:

Slice

```
exception PhoneEntryNotExists {
    string name;
    string phNum;
};

interface PhoneEntryFactory {
    PhoneEntry* create(string name, string phNum)
        throws PhoneEntryExists;
    void destroy(PhoneEntry* pe)      // Bad idea!
        throws PhoneEntryNotExists;
};
```

While this works (and certainly can be implemented without problems), it is generally a bad idea. For one, an immediate problem we need to deal with is what should happen if a client passes a proxy to an already-destroyed object to `destroy`. We could raise an `ObjectNotExistException` to indicate this, but that is not a good idea because it makes it ambiguous as to which object does not exist: the factory, or the entry. (By convention, if a client receives an `ObjectNotExistException` for an invocation, what does not exist is the object the operation was targeted at, not some other object that in turn might be contacted by the operation.) This forces us to add a separate `PhoneEntryNotExists` exception to deal with the error condition, which makes the interface a little more complex.

A second and more serious problem with this design is that, in order to destroy an entry, the client must not only know which entry to destroy, but must also know *which factory created the entry*. For our example, with only a single factory, this is not a real concern. However, for more complex systems with dozens of factories (possibly in multiple server processes), it rapidly becomes a problem: for each object, the application code somehow has to keep track of which factory created what object; if any part of the code ever loses track of where an object originally came from, it can no longer destroy that object.

Of course, we could mitigate the problem by adding an operation to the `PhoneEntry` interface that returns a proxy to its factory. That way, clients could ask each object to provide the factory that created the object. However, that needlessly complicates the Slice definitions and really is just a band-aid on a fundamentally flawed design. A much better choice is to add the `destroy` operation to the `PhoneEntry` interface instead:

Slice

```
interface PhoneEntry {
    idempotent string name();
    idempotent string getNumber();
    idempotent void setNumber(string phNum);
    void destroy();
};
```

With this approach, there is no need for clients to somehow keep track of which factory created what object. Instead, given a proxy to a `PhoneEntry` object, a client simply invokes the `destroy` operation on the object and the `PhoneEntry` obligingly commits suicide. Note that we also no longer need a separate exception to indicate the "object does not exist" condition because we can raise `ObjectNotExistException` instead — the exception exists precisely to indicate this condition and, because `destroy` is now an operation on the phone entry itself, there is no ambiguity about which object it is that does not exist.

Topics

Copyright 2016, ZeroC, Inc.

- Idempotency and Life Cycle Operations
- Implementing a destroy Operation
- Cleaning Up a Destroyed Servant
- Life Cycle and Collection Operations
- Life Cycle and Normal Operations

See Also

- Object Creation

Copyright 2016, ZeroC, Inc.

Idempotency and Life Cycle Operations

You may be tempted to write the life cycle operations as follows:

Slice

```
interface PhoneEntry {
    // ...
    idempotent void destroy(); // Wrong!

};

interface PhoneEntryFactory {
    idempotent PhoneEntry* create(string name, string phNum)
        throws PhoneEntryExists;
};
```

The idea is that `create` and `destroy` can be idempotent operations because it is safe to let the Ice run time retry the operation in the event of a temporary network failure. However, this assumption is not true. To see why, consider the following scenario:

1. A client invokes `destroy` on a phone entry.
2. The Ice run time sends the request to the server on the wire.
3. The connection goes down just after the request was sent, but before the reply for the request arrives in the client. It so happens that the request was received by the server and acted upon, and the reply from the server back to the client is lost because the connection is still down.
4. The Ice run time tries to read the reply for the request and realizes that the connection has gone down. Because the operation is marked idempotent, the run time attempts an automatic retry by re-establishing the connection and sending the request a second time, which happens to work.
5. The server receives the request to destroy the entry but, because the entry is destroyed already, the server returns an `ObjectNotExistException` to the client, which the Ice run time passes to the application code.
6. The application receives an `ObjectNotExistException` and falsely concludes that it tried to destroy a non-existent object when, in fact, the object did exist and was destroyed as intended.

A similar scenario can be constructed for `create`: in that case, the application will receive a `PhoneEntryExists` exception when, in fact, the entry did not exist and was created successfully.

These scenarios illustrate that `create` and `destroy` are *never* idempotent: sending one `create` or `destroy` invocation for a particular object is not the same as sending two invocations: the outcome depends on whether the first invocation succeeded or not, so `create` and `destroy` are not idempotent.

See Also

- Idempotent Operations
- Automatic Retries

Copyright 2016, ZeroC, Inc.

Implementing a destroy Operation

As far as the Ice run time is concerned, the act of destroying an Ice object is to remove the mapping between its proxy and its servant. In other words, an Ice object is destroyed when we remove its entry from the Active Servant Map (ASM). Once the ASM entry is gone, incoming operations for the object raise `ObjectNotExistException`, as they should.

On this page:

- Object Destruction and Concurrency
- Concurrent Execution of Life Cycle and Non-Life Cycle Operations

Object Destruction and Concurrency

Here is the simplest version of `destroy`:

C++

```
void
PhoneEntryI::destroy(const Current& c)
{
    try {
        c.adapter->remove(c.id);
    } catch (const Ice::NotRegisteredException&)
        throw Ice::ObjectNotExistException(__FILE__, __LINE__);
    }
}
```

The implementation removes the ASM entry for the servant, thereby destroying the Ice object. If the entry does not exist (presumably, because the object was destroyed previously), `destroy` throws an `ObjectNotExistException`, as you would expect.

The ASM entry is removed as soon as `destroy` calls `remove` on the object adapter. Assuming that we implement `create` as we saw earlier, so no other part of the code retains a smart pointer to the servant, this means that the ASM holds the only smart pointer to the servant, so the servant's reference count is 1.

We use smart pointers in C++, which are analogous to object references in languages such as Java and C#.

Once the ASM entry is removed (and its smart pointer destroyed), the reference count of the servant drops to zero. In C++, this triggers a call to the destructor of the servant, and the heap-allocated servant is deleted just as it should be; in languages such as Java and C#, this makes the servant eligible for garbage collection, so it will be deleted eventually as well.

Things get more interesting if we consider concurrent scenarios. One such scenario involves concurrent calls to `create` and `destroy`. Suppose we have the following sequence of events:

1. Client A creates a phone entry.
2. Client A passes the proxy for the entry to client B.
3. Client A destroys the entry.
4. Client A calls `create` for the same entry (passing the same name, which serves as the object identity) and, concurrently, client B calls `destroy` on the entry.

Clearly, something is strange about this scenario, because it involves two clients asking for conflicting things, with one client trying to create an object that existed previously, while another client tries to destroy the object that — unbeknownst to that client — was destroyed earlier.

Exactly what is seen by client A and client B depends on how the operations are dispatched in the server. In particular, the outcome depends on the order in which the calls on the object adapter to `add` (in `create`) and `remove` (in `destroy`) on the servant are executed:

- If the thread processing client A's invocation executes `add` before the thread processing client B's invocation, client A's call to `add` succeeds. Internally, the calls to `add` and `remove` are serialized, and client B's call to `remove` blocks until client A's call to `add` has completed. The net effect is that both clients see their respective invocations complete successfully.
- If the thread processing client B's invocation executes `remove` before the thread processing client A's invocation executes `add`, client B's thread receives a `NotRegisteredException`, which results in an `ObjectNotExistException` in client B. Client A's thread then successfully calls `add`, creating the object and returning its proxy.

Copyright 2016, ZeroC, Inc.

This example illustrates that, if life cycle operations interleave in this way, the outcome depends on thread scheduling. However, as far as the Ice run time is concerned, doing this is perfectly safe: concurrent access does not cause problems for memory management or the integrity of data structures.

The preceding scenario allows two clients to attempt to perform conflicting operations. This is possible because clients can control the object identity of each phone entry: if the object identity were hidden from clients and assigned by the server (the server could assign a UUID to each entry, for example), the above scenario would not be possible. We will return to a more detailed discussion of such object identity issues in Object Identity and Uniqueness.

Concurrent Execution of Life Cycle and Non-Life Cycle Operations

This section applies to C++ only.

Another scenario relates to concurrent execution of ordinary (non-life cycle) operations and `destroy`:

- Client A holds a proxy to an existing object and passes that proxy to client B.
- Client B calls the `setNumber` operation on the object.
- Client A calls `destroy` on the object *while Client B's call to `setNumber` is still executing*.

The immediate question is what this means with respect to memory management. In particular, client A's thread calls `remove` on the object adapter while client B's thread is still executing inside the object. If this call to `remove` were to delete the servant immediately, it would delete the servant while client B's thread is still executing inside the servant, with potentially disastrous results.

The answer is that this cannot happen. Whenever the Ice run time dispatches an incoming invocation to a servant, it increments the servant's reference count for the duration of the call, and decrements the reference count again once the call completes. Here is what happens to the servant's reference count for the preceding scenario:

1. Initially, the servant is idle, so its reference count is at least 1 because the ASM entry stores a smart pointer to the servant. (The remainder of these steps assumes that the ASM stores the *only* smart pointer to the servant, so the reference count is exactly 1.)
2. Client B's invocation of `setNumber` arrives and the Ice run time increments the reference count to 2 before dispatching the call.
3. While `setNumber` is still executing, client A's invocation of `destroy` arrives and the Ice run time increments the reference count to 3 before dispatching the call.
4. Client A's thread calls `remove` on the object adapter, which destroys the smart pointer in the ASM and so decrements the reference to 2.
5. Either `setNumber` or `destroy` may complete first. It does not matter which call completes — either way, the Ice run time decrements the reference count as the call completes, so after one of these calls completes, the reference count drops to 1.
6. Eventually, when the final call (`setNumber` or `destroy`) completes, the Ice run time decrements the reference count once again, which causes the count to drop to zero. In turn, this triggers the call to `delete` (which calls the servant's destructor).

The net effect is that, while operations are executing inside a servant, the servant's reference count is always greater than zero. As the invocations complete, the reference count drops until, eventually, it reaches zero. However, that can only happen once no operations are executing, that is, once the servant is idle. This means that the Ice run time guarantees that a servant's destructor runs only once the final operation invocation has drained out of the servant, so it is impossible to "pull memory out from underneath an executing invocation".

For garbage-collected languages, such as C# and Java, the language run time provides the same semantics: while the servant can be reached via any reference in the application or the Ice run time, the servant will not be reclaimed by the garbage collector.

See Also

- The Active Servant Map
- Object Identity and Uniqueness
- Servant Activation and Deactivation

Copyright 2016, ZeroC, Inc.

Cleaning Up a Destroyed Servant

Here is a very simple implementation of our `PhoneEntryI` servant. (Methods are inlined for convenience only. Also note that, for the time being, this code ignores concurrency issues, which we return to in Life Cycle and Normal Operations.)

C++

```
class PhoneEntryI : public PhoneEntry {
public:
    PhoneEntryI(const string& name, const string& phNum)
        : _name(name), _phNum(phNum)
    {
    }

    virtual string
    name(const Current&) {
        return _name;
    }

    virtual string
    getNumber(const Current&) {
        return _phNum;
    }

    virtual void
    setNumber(const string& phNum, const Current&) {
        _phNum = phNum;
    }

    virtual void
    destroy(const Current& c) {
        try {
           c.adapter->remove(c.id);
        } catch (const Ice::NotRegisteredException&)
           throw Ice::ObjectNotExistException(__FILE__, __LINE__);
        }
    }

private:
    const string _name;
    string _phNum;
};
```

With this servant, `destroy` does just the right thing: it calls `delete` on the servant once the servant is idle, which in turn calls the destructor, so the memory used by the `_name` and `_phNum` data members is reclaimed.

However, real servants are rarely this simple. In particular, destruction of an Ice object may involve non-trivial actions, such as flushing a file, committing a transaction, making a remote call on another object, or updating a hardware device. For example, instead of storing the details of a phone entry in member variables, the servant could be implemented to store the details in a file; in that case, destroying the Ice object would require closing the file. Seeing that the Ice run time calls the destructor of a servant only once the servant becomes idle, the destructor would appear to be an ideal place to perform such actions, for example:

Copyright 2016, ZeroC, Inc.

C++

```
class PhoneEntryI : public PhoneEntry {
public:
    // ...

    ~PhoneEntryI()
    {
        _myStream.close(); // Bad idea
    }

private:
    fstream _myStream;
};
```

The problem with this code is that it can fail, for example, if the file system is full and buffered data cannot be written to the file. Such clean-up failure is a general issue for non-trivial servants: for example, a transaction can fail to commit, a remote call can fail if the network goes down, or a hardware device can be temporarily unresponsive.

If we encounter such a failure, we have a serious problem: we cannot inform the client of the error because, as far as the client is concerned, the `destroy` call completed just fine. The client will therefore assume that the Ice object was correctly destroyed. However, the system is now in an inconsistent state: the Ice object was destroyed (because its ASM entry was removed), but the object's state still exists (possibly with incorrect values), which can cause errors later.

Another reason for avoiding such state clean-up in C++ destructors is that destructors cannot throw exceptions: if they do, and do so in the process of being called during unwinding of the stack due to some other exception, the program goes directly to `terminate` and does not pass "Go". (There are a few exotic cases in which it is possible to throw from a destructor and get away with it but, in general, is an excellent idea to maintain the no-throw guarantee for destructors.) So, if anything goes wrong during destruction, we are in a tight spot: we are forced to swallow any exception that might be encountered by the destructor, and the best we can do is log the error, but not report it to the client.

Finally, using destructors to clean up servant state does not port well to languages such as Java and C#. For these languages, similar considerations apply to error reporting from a finalizer and, with Java, finalizers may not run at all. Therefore, we recommend that you perform any clean-up actions in the body of `destroy` instead of delaying clean-up until the servant's destructor runs.

Note that the foregoing does *not* mean that you cannot reclaim servant resources in destructors; after all, that is what destructors are for. But it *does* mean that you should not try to reclaim resources from a destructor if the attempt can fail (such as deleting records in an external system as opposed to, for example, deallocating memory or adjusting the value of variables in your program).

See Also

- Life Cycle and Normal Operations

Copyright 2016, ZeroC, Inc.

Life Cycle and Collection Operations

On this page:

- Factories and Collection Operations
- Cyclic Dependencies between Factories and Objects

Factories and Collection Operations

The factory we defined earlier is what is known as a *pure* object factory because `create` is the only operation it provides. However, it is common for factories to do double duty and also act as collection managers that provide additional operations, such as `list` and `find`:

Slice

```
// ...

sequence<PhoneEntry*> PhoneEntries;

interface PhoneEntryFactory {
    PhoneEntry* create(string name, string phNum)
        throws PhoneEntryExists;

    idempotent PhoneEntry find(string name);
    idempotent PhoneEntries list();
};
```

`find` returns a proxy for the phone entry with the given name, and a null proxy if no such entry exists. `list` returns a sequence that contains the proxies of all existing entries.

Here is a simple implementation of `find`:

C++

```
PhoneEntryPrx
PhoneEntryFactory::find(const string& name, const Current& c)
{
    CommunicatorPtr comm = c.adapter->getCommunicator();

    PhoneEntryPrx pe;
    Identity id = comm->stringToIdentity(name);
    if (c.adapter->find(id)) {
        pe = PhoneEntryPrx::uncheckedCast(c.adapter->createProxy(id));
    }
    return pe;
}
```

If an entry exists in the Active Servant Map (ASM) for the given name, the code creates a proxy for the corresponding Ice object and returns it. This code works correctly even for threaded servers: because the look-up of the identity in the ASM is atomic, there is no problem with other threads concurrently modifying the ASM (for example, while servicing calls from other clients to `create` or `destroy`).

Cyclic Dependencies between Factories and Objects

Unfortunately, implementing `list` is not as simple because it needs to iterate over the collection of entries, but the object adapter does not

Copyright 2016, ZeroC, Inc.

provide any iterator for ASM entries.

The reason for this is that, during iteration, the ASM would have to be locked to protect it against concurrent access, but locking the ASM would prevent call dispatch during iteration and easily cause deadlocks.

Therefore, we must maintain our own list of entries inside the factory:

C++

```
class PhoneEntryFactoryI : public PhoneEntryFactory
{
public:
    // ...

    void remove(const string&, const ObjectAdapterPtr&) {
        IceUtil::Mutex::Lock lock(_namesMutex);

        set<string>::iterator i = _names.find(name);
        if (i != _names.end())
            _names.erase(i);
    }

private:
    IceUtil::Mutex _namesMutex;
    set<string> _names;
};
```

The idea is to have a set of names of existing entries, and to update that set in `create` and `destroy` as appropriate. However, for threaded servers, that raises a concurrency issue: if we have clients that can concurrently call `create`, `destroy`, and `list`, we need to interlock these operations to avoid corrupting the `_names` set (because STL containers are not thread-safe). This is the purpose of the mutex `_names Mutex` in the factory: `create`, `destroy`, and `list` can each lock this mutex to ensure exclusive access to the `_names` set.

Another issue is that our implementation of `destroy` must update the set of entries that is maintained by the factory. This is the purpose of the `remove` member function: it removes the specified name from the `_names` set (of course, under protection of the `_namesMutex` lock). However, `destroy` is a method on the `PhoneEntryI` servant, whereas `remove` is a method on the factory, so the servant must know how to reach the factory. Because the factory is a singleton, we can fix this by adding a static `_factory` member to the `PhoneEntryI` class:

C++

```
class PhoneEntryI : public PhoneEntry {
public:
    // ...

    static PhoneEntryFactoryIPtr _factory;

private:
    const string _name;
    string _phNum;
};
```

The code in `main` then creates the factory and initializes the static member variable, for example:

Copyright 2016, ZeroC, Inc.

C++

```
PersonI::_factory = new PersonFactoryI;

// Add factory to ASM and activate object
// adapter here...
```

This works, but it leaves a bad taste in our mouth because it sets up a cyclic dependency between the phone entry servants and the factory: the factory knows about the servants, and each servant knows about the factory so it can call `remove` on the factory. In general, such cyclic dependencies are a bad idea: if nothing else, they make a design harder to understand.

We could remove the cyclic dependency by moving the `_names` set and its associated mutex into a separate class instance that is referenced from both `PhoneEntryFactoryI` and `PhoneEntryI`. That would get rid of the cyclic dependency as far as the C++ type system is concerned but, as we will see later, it would not really help because the factory and its servants turn out to be mutually dependent regardless (because of concurrency issues). So, for the moment, we'll stay with this design and examine better alternatives after we have explored the concurrency issues in more detail.

With this design, we can implement `list` as follows:

C++

```
PhoneEntries
PhoneEntryFactoryI::list(const Current& c)
{
    Mutex::Lock lock(_namesMutex);

    CommunicatorPtr comm = c.adapter->getCommunicator();

    PhoneEntries pe;
    set<string>::const_iterator i;
    for (i = _names.begin(); i != _names.end(); ++i) {
        ObjectPrx o = c.adapter->createProxy(comm->stringToIdentity(name));
        pe.push_back(PhoneEntryPrx::uncheckedCast(o));
    }

    return pe;
}
```

Note that `list` acquires a lock on the mutex, to prevent concurrent modification of the `_names` set by `create` and `destroy`. In turn, our `create` implementation now also locks the mutex:

Copyright 2016, ZeroC, Inc.

```
PhoneEntryPrx
PhoneEntryFactory::create(const string& name, const string& phNum,
const Current& c)
{
    Mutex::Lock lock(_namesMutex);

    PhoneEntryPrx pe;
    try {
        CommunicatorPtr comm = c.adapter->getCommunicator();
        PhoneEntryPtr servant = new PhoneEntryI(name, phNum);
        pe = PhoneEntryPrx::uncheckedCast(c.adapter->add(servant,
                                          comm->stringToIdentity(name)));
    } catch (const Ice::AlreadyRegisteredException&) {
        throw PhoneEntryExists(name, phNum);
    }
    _names.insert(name);

    return pe;
}
```

With this implementation, we are safe if `create` and `list` run concurrently: only one of the two operations can acquire the lock at a time, so there is no danger of corrupting the `_names` set.

`destroy` is now trivial to implement: it simply removes the ASM entry and calls `remove` on the factory:

C++

```
void
PhoneEntryI::destroy(const Current& c)
{
    // Note: not quite correct yet.
    c.adapter->remove(_name);
    _factory->remove(_name, c.adapter);
}
```

Note that this is not quite correct because we have not yet considered concurrency issues for destroy. We will consider this issue next.

See Also

- Object Creation
- The Active Servant Map

Copyright 2016, ZeroC, Inc.

Life Cycle and Normal Operations

So far, we have mostly ignored the implementations of the `getNumber` and `setNumber` operations we defined in our discussion of object creation. Obviously, `getNumber` and `setNumber` must be interlocked against concurrent access — without this interlock, concurrent requests from clients could result in one thread writing to the `_phNum` member while another thread is reading it, with unpredictable results. (Conversely, the `name` operation need not have an interlock because the name of a phone entry is immutable.) To interlock `getNumber` and `setNumber`, we can add a mutex `_m` to `PhoneEntryI`:

C++

```
class PhoneEntryI : public PhoneEntry {
public:
    // ...

    static PhoneEntryFactoryIPtr _factory;

private:
    const string _name;
    string _phNum;
    Mutex  _m;
};
```

The `getNumber` and `setNumber` implementations then lock `_m` to protect `_phNum` from concurrent access:

Copyright 2016, ZeroC, Inc.

C++

```
string
PhoneEntryI::name(const Current&)
{
    return _name;
}

string
PhoneEntryI::getNumber(const Current&)
{
    // Incorrect implementation!

    Mutex::Lock lock(_m);

    return _phNum;
}

void
PhoneEntryI::setNumber(const string& phNum, const Current&)
{
    // Incorrect implementation!

    Mutex::Lock lock(_m);

    _phNum = phNum;
}
```

This looks good but, as it turns out, `destroy` throws a spanner in the works: as shown, this code suffers from a rare, but real, race condition. Consider the situation where a client calls `destroy` at the same time as another client calls `setNumber`. In a server with a thread pool with more than one thread, the calls can be dispatched in separate threads and can therefore execute concurrently.

The following sequence of events can occur:

- The thread dispatching the `setNumber` call locates the servant, enters the operation implementation, and is suspended by the scheduler immediately on entry to the operation, before it can lock `_m`.
- The thread dispatching the `destroy` call locates the servant, enters `destroy`, successfully removes the servant from the Active Servant Map (ASM) and the `_names` set, and returns.
- The thread that was suspended in `setNumber` is scheduled again, locks `_m`, and now operates on a conceptually already-destroyed Ice object.

Note that, even if we lock `_m` in `destroy`, the problem persists:

Copyright 2016, ZeroC, Inc.

C++

```
void
PhoneEntryI::destroy(const Current& c)
{
    // Note: still not correct.

    IceUtil::Mutex::Lock lock(_m);

    c.adapter->remove(_name);
    _factory->remove(_name, c.adapter);
}
```

Even though `destroy` now locks `_m` and so cannot run concurrently with `getNumber` and `setNumber`, the preceding scenario can still arise. The problem here is that a thread can enter the servant and be suspended before it gets a chance to acquire a lock. With the code as it stands, this is not a problem: `setNumber` will simply update the `_phNum` member variable in a servant that no longer has an ASM entry. In other words, the Ice object is already destroyed — it just so happens that the servant for that Ice object is still hanging around because there is still an operation executing inside it. Any updates to the servant will succeed (even though they are useless because the servant's destruct or will run as soon as the last invocation leaves the servant.)

Note that this scenario is not unique to C++ and can arise even with Java synchronized operations: in that case, a thread can be suspended just after the Ice run time has identified the target servant, but before it actually calls the operation on the target servant. While the thread is suspended, another thread can execute `destroy`.

While this race condition does not affect our implementation, it does affect more complex applications, particularly if the servant modifies external state, such as a file system or database. For example, `setNumber` could modify a file in the file system; in that case, `destroy` woul d delete that file and probably close a file descriptor or stream. If we were to allow `setNumber` to continue executing after `destroy` has already done its job, we would likely encounter problems: `setNumber` might not find the file where it expects it to be or try to use the closed file descriptor and return an error; or worse, `setNumber` might end up re-creating the file in the process of updating the already-destroyed entry's phone number. (What exactly happens depends on how we write the code for each operation.)

Of course, we can try to anticipate these scenarios and handle the error conditions appropriately, but doing this for complex systems with complex servants rapidly gets out of hand: in each operation, we would have to ask ourselves what might happen if the servant is destroyed concurrently and, if so, take appropriate recovery action.

It is preferable to instead deal with interleaved invocations of `destroy` and other operations in a systematic fashion. We can do this by adding a `_destroyed` member to the `PhoneEntryI` servant. This member is initialized to false by the constructor and set to true by `destr oy`. On entry to every operation (including `destroy`), we lock the mutex, test the `_destroyed` flag, and throw `ObjectNotExistExceptio n` if the flag is set:

C++

```
class PhoneEntryI : public PhoneEntry {
public:
    // ...

    static PhoneEntryFactoryIPtr _factory;

private:
    const string _name;
    string _phNum;
    bool   _destroyed;
    Mutex  _m;
};
```


Copyright 2016, ZeroC, Inc.

```
PhoneEntryI::PhoneEntryI(const string& name, const string& phNum)
    : _name(name), _phNum(phNum), _destroyed(false)
{
}

string
PhoneEntryI::name(const Current&)
{
    Mutex::Lock lock(_m);

    if (_destroyed)
        throw ObjectNotExistException(__FILE__, __LINE__);

    return _name;
}

string
PhoneEntryI::getNumber(const Current&)
{
    Mutex::Lock lock(_m);

    if (_destroyed)
        throw ObjectNotExistException(__FILE__, __LINE__);

    return _phNum;
}

void
PhoneEntryI::setNumber(const string& phNum, const Current&)
{
    Mutex::Lock lock(_m);

    if (_destroyed)
        throw ObjectNotExistException(__FILE__, __LINE__);

    _phNum = phNum;
}

void
PhoneEntryI::destroy(const Current& c)
{
    Mutex::Lock lock(_m);

    if (_destroyed)
        throw ObjectNotExistException(__FILE__, __LINE__);

    _destroyed = true;
    c.adapter->remove(_name);
```


Copyright 2016, ZeroC, Inc.

```
    _factory->remove(_name, c.adapter); // Dubious!
}
```

If you are concerned about the repeated code on entry to every operation, you can put that code into a member function or base class to make it reusable (although the benefits of doing so are probably too minor to make this worthwhile).

Using the `_destroyed` flag, if an operation is dispatched and suspended before it can lock the mutex and, meanwhile, `destroy` runs to completion in another thread, it becomes impossible for an operation to operate on the state of such a "zombie" servant: the test on entry to each operation ensures that any operation that runs after `destroy` immediately raises `ObjectNotExistException`.

Also note the "dubious" comment in `destroy`: the operation first locks `_m` and, while holding that lock, calls `remove` on the factory, which in turn locks its own `_namesMutex`. This is not wrong as such, but as we will see shortly, it can easily lead to deadlocks if we modify the application later.

See Also

- Object Creation
- The Active Servant Map
- Implementing a destroy Operation
- The Ice Threading Model

Copyright 2016, ZeroC, Inc.

Removing Cyclic Dependencies

We mentioned earlier that factoring the `_names` set and its mutex into a separate class instance does not really solve the cyclic dependency problem, at least not in general. To see why, suppose that we want to extend our factory with a new `getDetails` operation:

Slice

```
// ...

struct Details {
    PhoneEntry* proxy;
    string name;
    string phNum;
};

sequence<Details> DetailsSeq;

interface PhoneEntryFactory {
    // ...

    DetailsSeq getDetails();
};
```

This type of operation is common in collection managers: instead of returning a simple list of proxies, `getDetails` returns a sequence of structures, each of which contains not only the object's proxy, but also some of the state of the corresponding object. The motivation for this is performance: with a plain list of proxies, the client, once it has obtained the list, is likely to immediately follow up with one or more remote calls for each object in the list in order to retrieve their state (for example, to display the list of objects to the user). Making all these additional remote procedure calls is inefficient, and an operation such as `getDetails` gets the job done with a single RPC instead.

To implement `getDetails` in the factory, we need to iterate over the set of entries and invoke the `getNumber` operation on each object. (These calls are collocated and therefore very efficient, so they do not suffer the performance problem that a client calling the same operations would suffer.) However, this is potentially dangerous because the following sequence of events is possible:

- Client A calls `getDetails`.
- The implementation of `getDetails` must lock `_namesMutex` to prevent concurrent modification of the `_names` set during iteration.
- Client B calls `destroy` on a phone entry.
- The implementation of `destroy` locks the entry's mutex `_m`, sets the `_destroyed` flag, and then calls `remove`, which attempts to lock `_namesMutex` in the factory. However, `_namesMutex` is already locked by `getDetails`, so `remove` blocks until `_m` is unlocked again.
- `getDetails`, while iterating over its set of entries, happens to call `getNumber` on the entry that is currently being destroyed by client B. `getNumber`, in turn, tries to lock its mutex `_m`, which is already locked by `destroy`.

At this point, the server deadlocks: `getDetails` holds a lock on `_namesMutex` and waits for `_m` to become available, and `destroy` holds a lock on `_m` and waits for `_namesMutex` to become available, so neither thread can make progress.

To get rid of the deadlock, we have two options:

- Rearrange the locking such that deadlock becomes impossible.
- Abandon the idea of calling back from the servants into the factory and use *reaping* instead.

We will explore both options in the following pages.

Topics

- Acquiring Locks without Deadlocks
- Reaping Objects

See Also

- Life Cycle and Collection Operations

Copyright 2016, ZeroC, Inc.

Acquiring Locks without Deadlocks

For our example, it is fairly easy to avoid the deadlock caused by cyclic dependencies: instead of holding the lock for the duration of `destroy`, we set the `_destroyed` flag under protection of the lock and unlock `_m` again before calling `remove` on the factory:

C++

```
void
PhoneEntryI::destroy(const Current& c)
{
    {
        Mutex::Lock lock(_m);

        if (_destroyed)
            throw ObjectNotExistException(__FILE__, __LINE__);

        _destroyed = true;

    } // _m is unlocked here.

    _factory->remove(_name, c.adapter);
}
```

Now deadlock is impossible because no function holds more than one lock, and no function calls another function while it holds a lock. However, rearranging locks in this fashion can be quite difficult for complex applications. In particular, if an application uses callbacks that do complex things involving several objects, it can be next to impossible to prove that the code is free of deadlocks. The same is true for applications that use condition variables and suspend threads until a condition becomes true.

At the core of the problem is that concurrency can create circular locking dependencies: an operation on the factory (such as `getDetails`) can require the same locks as a concurrent call to `destroy`. This is one reason why threaded code is harder to write than sequential code — the interactions among operations require locks, but dependencies among these locks are not obvious. In effect, locks set up an entirely separate and largely invisible set of dependencies. For example, it was easy to spot the mutual dependency between the factory and the servants due to the presence of `remove`; in contrast, it was much harder to spot the lurking deadlock in `destroy`. Worse, deadlocks may not be found during testing and discovered only after deployment, when it is much more expensive to rectify the problem.
See Also

- Removing Cyclic Dependencies

Copyright 2016, ZeroC, Inc.

Reaping Objects

Instead of trying to arrange code such that it is deadlock-free in the presence of callbacks and cyclic dependencies, it is often easier to change the code to avoid the callbacks entirely and to use an approach known as *reaping*:

- `destroy` marks the servant as destroyed and removes the Active Servant Map (ASM) entry as usual, but it does not call back into the factory to update the `_names` set.
- Whenever a collection manager operation, such as `list`, `getDetails`, or `find` is called, the factory checks for destroyed servants and, if it finds any, removes them from the `_names` set.

Reaping can make for a much cleaner design because it avoids both the cyclic type dependency and the cyclic locking dependency.

On this page:

- A Simple Reaping Implementation
- Alternative Reaping Implementations

A Simple Reaping Implementation

To implement reaping, we need to change our `PhoneEntryI` definition a little. It no longer has a static `_factory` smart pointer back to the factory (because it no longer calls `remove`). Instead, the servant now provides a member function `_isZombie` that the factory calls to check whether the servant was destroyed some time in the past:

C++

```
class PhoneEntryI : public PhoneEntry {
public:
    // ...

    bool _isZombie() const;

private:
    const string _name;
    string _phNum;
    bool   _destroyed;
    Mutex  _m;
};
```

The implementation of `_isZombie` is trivial: it returns the `_destroyed` flag under protection of the lock:

C++

```
bool
PhoneEntryI::_isZombie() const
{
    Mutex::Lock lock(_m);

    return _destroyed;
}
```

The `destroy` operation no longer calls back into the factory to update the `_names` set; instead, it simply sets the `_destroyed` flag and removes the ASM entry:

Copyright 2016, ZeroC, Inc.

C++

```
void
PhoneEntryI::destroy(const Current& c)
{
    Mutex::Lock lock(_m);

    if (_destroyed)
        throw ObjectNotExistException(__FILE__, __LINE__);

    _destroyed = true;
    c.adapter->remove(c.id);
}
```

The factory now, instead of storing just the names of existing servants, maintains a map that maps the name of each servant to its smart pointer:

C++

```
class PhoneEntryFactoryI : public PhoneEntryFactory
{
public:
    // Constructor and Slice operations here...

private:
    typedef map<string, PhoneEntryIPtr> PMap;
    PMap _entries;
    Mutex _entriesMutex;
};
```

During `create` (and other operations, such as `list`, `getDetails`, and `find`), we scan for zombie servants and remove them from the `_entries` map:

Copyright 2016, ZeroC, Inc.

C++

```
PhoneEntryPrx
PhoneEntryFactory::create(const string& name, const string& phNum,
const Current& c)
{
    Mutex::Lock lock(_entriesMutex);

    PhoneEntryPrx pe;
    PhoneEntryIPtr servant = new PhoneEntryI(name, phNum);

    // Try to create new object.
    //
    try {
        CommunicatorPtr comm = c.adapter->getCommunicator();
        pe = PhoneEntryPrx::uncheckedCast(
            c.adapter->add(servant, comm->stringToIdentity(name)));
    } catch (const Ice::AlreadyRegisteredException&) {
        throw PhoneEntryExists(name, phNum);
    }

    // Scan for zombies.
    //
    PMap::iterator i = _entries.begin();
    while (i != _entries.end())
    {
        if (i->second->_isZombie())
            _entries.erase(i++);
        else
            ++i;
    }
    _entries[name] = servant;

    return pe;
}
```

The implementations of `list`, `getDetails`, and `find` scan for zombies as well. Because they need to iterate over the existing entries anyway, reaping incurs essentially no extra cost:

Copyright 2016, ZeroC, Inc.

C++

```
PhoneEntries
PhoneEntryFactoryI::list(const Current& c)
{
    Mutex::Lock lock(_entriesMutex);

    CommunicatorPtr comm = c.adapter->getCommunicator();

    PhoneEntries pe;
    PMap::iterator i = _entries.begin();
    for (i = _entries.begin(); i != _entries.end(); ++i) {
        if (i->second->_isZombie()) {
            _entries.erase(i++);
        } else {
            ObjectPrx o = c.adapter->createProxy(comm->stringToIdentity(i->first));
            pe.push_back(PhoneEntryPrx::uncheckedCast(o));
            ++i;
        }
    }

    return pe;
}

// Similar for getDetails and find...
```

This is a much cleaner design: there is no cyclic dependency between the factory and the servants, either implicit (in the type system) or explicit (as a locking dependency). Moreover, the implementation is easier to understand once you get used to the idea of reaping: there is no need to follow complex callbacks and to carefully analyze the order of lock acquisition. (Note that, depending on how state is maintained for servants, you may also need to reap during start-up and shutdown.)

In general, we recommend that you use a reaping approach in preference to callbacks for all but the most trivial applications: it simply is a better approach that is easier to maintain and understand.

Alternative Reaping Implementations

You may be concerned that reaping increases the cost of `create` from *O*(log *n*) to *O*(*n*) because `create` now iterates over all existing entries and locks and unlocks a mutex in each servant (whereas, previously, it simply added each new servant to the `_names` set). Often, this is not an issue because life cycle operations are called infrequently compared to normal operations. However, you will notice the additional cost if you have a large number of servants (in the thousands or more) and life cycle operations are called frequently.

If you find that `create` is a bottleneck (by profiling, not by guessing!), you can change to a more efficient implementation by adding zombie servants to a separate zombie list. Reaping then iterates over the zombie list and removes each servant in the zombie list from the `_entries` map before clearing the zombie list. This reduces the cost of reaping to be proportional to the number of zombie servants instead of the total number of servants. In addition, it allows us to remove the `_isZombie` member function and to lock and unlock `_entriesMutex` only once instead of locking a mutex in each servant as part of `_isZombie`. We will see such an implementation when we revisit the file system application.

You may also be concerned about the number of zombie servants that can accumulate in the server if `create` is not called for some time. For most applications, this is not a problem: the servants occupy memory, but no other resources because `destroy` can clean up scarce resources, such as file descriptors or network connections before it turns the servant into a zombie. If you really need to prevent accumulation of zombie servants, you can reap from a background thread that runs periodically, or you can count the number of zombies and trigger a reaping pass once that number exceeds some threshold.

See Also

- Removing Cyclic Dependencies
- The Active Servant Map

Copyright 2016, ZeroC, Inc.

- Object Life Cycle for the File System Application

Copyright 2016, ZeroC, Inc.

Object Identity and Uniqueness

When creating an object with a factory, it may not be a good idea to allow clients to control the object identity of the Ice objects they create. Here is the scenario from our previous discussion, re-cast in terms of our phone book application:

1. Client A creates a new phone entry for Fred.
2. Client A passes the proxy for the Fred entry as a parameter of a remote call to another part of the system, say, server B.
3. Server B remembers Fred's proxy.
4. Client A decides that the entry for Fred is no longer needed and calls Fred's destroy operation.
5. Some time later, client C creates a new phone entry for a different person whose name also happens to be Fred.
6. Server B decides to get Fred's phone number by calling `getNumber` on the proxy it originally obtained from client A.

At this point, things are likely to go wrong: server B thinks that it has obtained the phone number of the original Fred, but that entry no longer exists and has since been replaced by a new entry for a different person (who presumably has a different phone number).

What has happened here is that Fred has been reincarnated because the same object identity was used for two different objects. In general, such reused object identities are a bad idea. For example, consider the following interfaces:

Slice

```
interface Process {
    void launch(); // Start process.
};

interface Missile {
    void launch(); // Kill lots of people.
};
```

Replaying the preceding scenario, if client A creates a `Process` object called "Thunderbird" and destroys that object again, and client C creates a `Missile` object called "Thunderbird", when server B calls `launch`, it will launch a missile instead of a process.

To be fair, in reality, this scenario is unlikely because it tacitly assumes that both objects are implemented by the same object adapter but, in a realistic scenario, the same server would be unlikely to implement both launching of processes and missiles. However, if you have objects that derive from common base interfaces, so objects of different types share the same operation names, this problem is real: operation invocations can easily end up in a destroyed and later recreated object.

Specifically, the preceding scenario illustrates that, when the Ice run time dispatches a request, exactly three items determine where the request ends up being processed:

- the endpoint at which the server listens for incoming requests
- the identity of the Ice object that is the target of the request
- the name of the operation that is to be invoked on the Ice object

If object identities are insufficiently unique, a request intended for one object can end up being sent to a completely different object, provided that the original object used the same identity, that both provide an operation with the same name, and that the parameters passed to one operation happen to decode correctly when interpreted as the parameters to the other operation. (This is rare, but not impossible, depending on the type and number of parameters.)

The crucial question is, what do we mean by "insufficiently unique"? As far as the call dispatch is concerned, identities must be unique only per object adapter. This is because the ASM does not allow you to add two entries with the same object identity; by enforcing this, the Active Servant Map (ASM) ensures that each object identity belongs to exactly one servant. (Note that the converse, namely, that servants in ASM entries must be unique, is *not* the case: the ASM allows you to map different object identities to the same servant, which is useful to, for example, implement stateless facade objects — default servants are used for this purpose.) So, as far as the Ice run time is concerned, it is perfectly OK to reuse object identities for different Ice objects.

Note that the Ice run time cannot prevent reuse of object identities either. Doing so would require the run time to remember every object identity that has ever been used, which does not scale. Instead, the Ice run time makes the application responsible for ensuring that object identities are "sufficiently unique".

You can deal with the identity reuse problem in several ways. One option is to do nothing and simply ignore the problem. While this sounds facetious, it is a viable option for many applications because, due to their nature, identity reuse is simply impossible. For example, if you use a social security number as a person's identity, the problem cannot arise because the social security number of a deceased person is not given to another person.

Copyright 2016, ZeroC, Inc.

Another option is to allow identity reuse and to write your application such that it can deal with such identities: if nothing bad happens when an identity is reused, there is no problem. (This is the case if you know that the life cycles of the proxies for two different objects with the same identity can never overlap.)

The third option is to ensure that object identities are guaranteed unique, for example, by establishing naming conventions that make reuse impossible, or by using a UUID as the object identity. This can be useful even for applications for which identity reuse does not pose a problem. For example, if you use well-known objects, globally-unique object identities allow you to move a server to a different machine without invalidating proxies to these objects that are held by clients.

In general, we recommend that if an Ice object naturally contains a unique item of state (such as a social security number), you should use that item as the object identity. On the other hand, if the natural object identity is insufficiently unique (as is the case with names of phone book entries), you should use a UUID as the identity. (This is particularly useful for anonymous transient objects, such as session objects, that may not have a natural identity.)

See Also

- Object Creation
- Object Identity
- The Active Servant Map
- Default Servants
- Servant Activation and Deactivation
- Understanding Object Life Cycle
- Well-Known Objects

Copyright 2016, ZeroC, Inc.

Object Life Cycle for the File System Application

Now that we have had a look at the issues around object life cycle, let us return to our file system application and add life cycle operations to it, so clients can create and destroy files and directories.

To destroy a file or directory, the obvious choice is to add a `destroy` operation to the `Node` interface:

Slice

```
module Filesystem {

    exception GenericError {
        string reason;
    };
    exception PermissionDenied extends GenericError {};
    exception NameInUse extends GenericError {};
    exception NoSuchName extends GenericError {};

    interface Node {
        idempotent string name();
        void destroy() throws PermissionDenied;
    };

    // ...
};
```

Note that `destroy` can throw a `PermissionDenied` exception. This is necessary because we must prevent attempts to destroy the root directory.

The `File` interface is unchanged:

Slice

```
module Filesystem {
    // ...

    sequence<string> Lines;

    interface File extends Node {
        idempotent Lines read();
        idempotent void write(Lines text) throws GenericError;
    };

    // ...
};
```

Note that, because `File` derives from `Node`, it inherits the `destroy` operation we defined for `Node`.

The `Directory` interface now looks somewhat different from the previous version:

- The `list` operation returns a sequence of structures instead of a list of proxies: for each entry in a directory, the `NodeDesc` structure provides the name, type, and proxy of the corresponding file or directory.
- Directories provide a `find` operation that returns the description of the nominated node. If the nominated node does not exist, the

Copyright 2016, ZeroC, Inc.

operation throws a `NoSuchName` exception.
- The `createFile` and `createDirectory` operations create a file and directory, respectively. If a file or directory already exists, the operations throw a `NameInUse` exception.

Here are the corresponding definitions:

Slice

```
module Filesystem {
    // ...

    enum NodeType { DirType, FileType };

    struct NodeDesc {
        string name;
        NodeType type;
        Node* proxy;
    };

    sequence<NodeDesc> NodeDescSeq;

    interface Directory extends Node {
        idempotent NodeDescSeq list();
        idempotent NodeDesc find(string name) throws NoSuchName;
        File* createFile(string name) throws NameInUse;
        Directory* createDirectory(string name) throws NameInUse;
    };
};
```

Note that this design is somewhat different from the factory we designed for the phone book application. In particular, we do not have a single object factory; instead, we have as many factories as there are directories, that is, each directory creates files and directories only in that directory.

The motivation for this design is twofold:

- Because all files and directories that can be created are immediate descendants of their parent directory, we avoid the complexities of parsing path names for a separator such as "/". This keeps our example code to manageable size. (A real-world implementation of a distributed file system would, of course, be able to deal with path names.)
- Having more than one object factory presents interesting implementation issues that we will explore in the following discussion.

Let's move on to the implementation of this design in C++ and Java. You can find the full code of the implementation (including languages other than C++ and Java) in the `Manual/lifecycle` directory of your Ice distribution.

Topics

- Implementing Object Life Cycle in C++
- Implementing Object Life Cycle in Java

See Also

- Slice for a Simple File System
- Object Creation

Copyright 2016, ZeroC, Inc.

Implementing Object Life Cycle in C++

The implementation of our life cycle design has the following characteristics:

- It uses UUIDs as the object identities for nodes to avoid object reincarnation problems.
- When `destroy` is called on a node, the node needs to destroy itself and inform its parent directory that it has been destroyed (because the parent directory is the node's factory and also acts as a collection manager for child nodes).

Note that, in contrast to the initial version, the entire implementation resides in a `FilesystemI` namespace instead of being part of the `Filesystem` namespace. Doing this is not essential, but is a little cleaner because it keeps the implementation in a namespace that is separate from the Slice-generated namespace.

On this page:

- Object Life Cycle Changes for the NodeI Class in C++
- Object Life Cycle Changes for the DirectoryI Class in C++
- Object Life Cycle Changes for the FileI Class in C++
- Object Life Cycle Concurrency Issues in C++

Object Life Cycle Changes for the `NodeI` Class in C++

To begin with, let us look at the definition of the `NodeI` class:

C++

```
namespace FilesystemI {

    class DirectoryI;
    typedef IceUtil::Handle<DirectoryI> DirectoryIPtr;

    class NodeI : public virtual Filesystem::Node {
    public:
        virtual std::string name(const Ice::Current&);
        Ice::Identity id() const;

    protected:
        NodeI(const std::string& name, const DirectoryIPtr& parent);

        const std::string _name;
        const DirectoryIPtr _parent;
        bool _destroyed;
        Ice::Identity _id;
        IceUtil::Mutex _m;
    };

    // ...
}
```

The purpose of the `NodeI` class is to provide the data and implementation that are common to both `FileI` and `DirectoryI`, which use implementation inheritance from `NodeI`.

As in the initial version, `NodeI` provides the implementation of the `name` operation and stores the name of the node and its parent directory in the `_name` and `_parent` members. (The root directory's `_parent` member is null.) These members are immutable and initialized by the constructor and, therefore, `const`.

The `_destroyed` member, protected by the mutex `_m`, prevents the race condition we discussed earlier. The constructor initializes `_destroyed` to `false` and creates an identity for the node (stored in the `_id` member):

Copyright 2016, ZeroC, Inc.

C++

```
FilesystemI::NodeI::NodeI(const string& name, const DirectoryIPtr& parent)
    : _name(name), _parent(parent), _destroyed(false)
{
    _id.name = parent ? IceUtil::generateUUID() : "RootDir";
}
```

The `id` member function returns a node's identity, stored in the `_id` data member. The node must remember this identity because it is a UUID and is needed when we create a proxy to the node:

C++

```
Identity
FilesystemI::NodeI::id() const
{
    return _id;
}
```

The data members of `NodeI` are protected instead of private to keep them accessible to the derived `FileI` and `DirectoryI` classes. (Because the implementation of `NodeI` and its derived classes is quite tightly coupled, there is little point in making these members private and providing separate accessors and mutators for them.)

The implementation of the Slice `name` operation simply returns the name of the node, but also checks whether the node has been destroyed:

C++

```
string
FilesystemI::NodeI::name(const Current&)
{
    IceUtil::Mutex::Lock lock(_m);

    if (_destroyed)
        throw ObjectNotExistException(__FILE__, __LINE__);

    return _name;
}
```

This completes the implementation of the `NodeI` base class.

Object Life Cycle Changes for the `DirectoryI` Class in C++

Next, we need to look at the implementation of directories. The `DirectoryI` class derives from `NodeI` and the Slice-generated `Directory` skeleton class. Of course, it must implement the pure virtual member functions for its Slice operations, which leads to the following (not yet complete) definition:

Copyright 2016, ZeroC, Inc.

C++

```
namespace FilesystemI {

    // ...

    class DirectoryI : virtual public NodeI,
                       virtual public Filesystem::Directory {
    public:
        virtual Filesystem::NodeDescSeq list(const Ice::Current&);
        virtual Filesystem::NodeDesc find(const std::string&, const Ice::Current&);
        Filesystem::FilePrx createFile(const std::string&, const Ice::Current&);
        Filesystem::DirectoryPrx createDirectory(const std::string&,
const Ice::Current&);
        virtual void destroy(const Ice::Current&);
        // ...

    private:
        // ...
    };
}
```

Each directory stores its contents in a map that maps the name of a directory to its servant:

C++

```
namespace FilesystemI {

    // ...

    class DirectoryI : virtual public NodeI,
                       virtual public Filesystem::Directory {
    public:
        // ...

        DirectoryI(const ObjectAdapterPtr& a,
                   const std::string& name,
                   const DirectoryIPtr& parent = 0);

        void removeEntry(const std::string& name);

    private:
        typedef std::map<std::string, NodeIPtr> Contents;
        Contents _contents;
        // ...
    };
}
```

Note that we use the inherited member `_m` to interlock operations.

The constructor simply initializes the `NodeI` base class:

Copyright 2016, ZeroC, Inc.

C++

```
FilesystemI::DirectoryI::DirectoryI(const string& name,
const DirectoryIPtr& parent)
    : NodeI(name, parent)
{
}
```

The `removeEntry` member function is called by the child to remove itself from its parent's `_contents` map:

C++

```
void
FilesystemI::DirectoryI::removeEntry(const string& name)
{
    IceUtil::Mutex::Lock lock(_m);
    Contents::iterator i = _contents.find(name);
    if(i != _contents.end())
    {
        _contents.erase(i);
    }
}
```

Here is the `destroy` member function for directories:

C++

```
void
FilesystemI::DirectoryI::destroy(const Current& c)
{
    if (!_parent)
        throw PermissionDenied("Cannot destroy root directory");
    {
        IceUtil::Mutex::Lock lock(_m);

        if (_destroyed)
            throw ObjectNotExistException(__FILE__, __LINE__);

        if (!_contents.empty())
            throw PermissionDenied("Cannot destroy non-empty directory");

        c.adapter->remove(id());
        _destroyed = true;
    }

    _parent->removeEntry(_name);
}
```

The code first prevents destruction of the root directory and then checks whether this directory was destroyed previously. It then acquires the

Copyright 2016, ZeroC, Inc.

lock and checks that the directory is empty. Finally, `destroy` removes the Active Servant Map (ASM) entry for the destroyed directory and removes itself from its parent's `_contents` map. Note that we call `removeEntry` outside the synchronization to avoid deadlocks.

The `createDirectory` implementation locks the mutex before checking whether the directory already contains a node with the given name (or an invalid empty name). If not, it creates a new servant, adds it to the ASM and the `_contents` map, and returns its proxy:

C++

```
DirectoryPrx
FilesystemI::DirectoryI::createDirectory(const string& name, const Current& c)
{
    IceUtil::Mutex::Lock lock(_m);

    if (_destroyed)
        throw ObjectNotExistException(__FILE__, __LINE__);

    if (name.empty() || _contents.find(name) != _contents.end())
        throw NameInUse(name);

    DirectoryIPtr d = new DirectoryI(name, this);
    ObjectPrx node = c.adapter->add(d, d->id());
    _contents[name] = d;
    return DirectoryPrx::uncheckedCast(node);
}
```

The `createFile` implementation is identical, except that it creates a file instead of a directory:

C++

```
FilePrx
FilesystemI::DirectoryI::createFile(const string& name, const Current& c)
{
    IceUtil::Mutex::Lock lock(_m);

    if (_destroyed)
        throw ObjectNotExistException(__FILE__, __LINE__);

    if (name.empty() || _contents.find(name) != _contents.end())
        throw NameInUse(name);

    FileIPtr f = new FileI(name, this);
    ObjectPrx node = c.adapter->add(f, f->id());
    _contents[name] = f;
    return FilePrx::uncheckedCast(node);
}
```

Here is the implementation of `list`:

Copyright 2016, ZeroC, Inc.

C++

```
NodeDescSeq
FilesystemI::DirectoryI::list(const Current& c)
{
    IceUtil::Mutex::Lock lock(_m);

    if (_destroyed)
        throw ObjectNotExistException(__FILE__, __LINE__);

    NodeDescSeq ret;
    for (Contents::const_iterator i = _contents.begin(); i != _contents.end(); ++i)
    {
        NodeDesc d;
        d.name = i->first;
        d.type = FilePtr::dynamicCast(i->second) ? FileType : DirType;
        d.proxy = NodePrx::uncheckedCast(c.adapter->createProxy(i->second->id()));
        ret.push_back(d);
    }
    return ret;
}
```

After acquiring the lock, the code iterates over the directory's contents and adds a `NodeDesc` structure for each entry to the returned vector.

The `find` operation proceeds along similar lines:

C++

```
NodeDesc
FilesystemI::DirectoryI::find(const string& name, const Current& c)
{
    IceUtil::Mutex::Lock lock(_m);

    if (_destroyed)
        throw ObjectNotExistException(__FILE__, __LINE__);

    Contents::const_iterator pos = _contents.find(name);
    if (pos == _contents.end())
        throw NoSuchName(name);

    NodeIPtr p = pos->second;
    NodeDesc d;
    d.name = name;
    d.type = FilePtr::dynamicCast(p) ? FileType : DirType;
    d.proxy = NodePrx::uncheckedCast(c.adapter->createProxy(p->id()));
    return d;
}
```

Object Life Cycle Changes for the `FileI` Class in C++

The constructor of `FileI` is trivial: it simply initializes the data members of its base class::

Copyright 2016, ZeroC, Inc.

C++

```
FilesystemI::FileI::FileI(const string& name, const DirectoryIPtr& parent)
    : NodeI(name, parent)
{
}
```

The implementation of the three member functions of the `FileI` class is also trivial, so we present all three member functions here:

C++

```
Lines
FilesystemI::FileI::read(const Current&)
{
    IceUtil::Mutex::Lock lock(_m);

    if (_destroyed)
        throw ObjectNotExistException(__FILE__, __LINE__);

    return _lines;
}

// Slice File::write() operation.

void
FilesystemI::FileI::write(const Lines& text, const Current&)
{
    IceUtil::Mutex::Lock lock(_m);

    if (_destroyed)
        throw ObjectNotExistException(__FILE__, __LINE__);

    _lines = text;
}

void
FilesystemI::FileI::destroy(const Current& c)
{
    {
        IceUtil::Mutex::Lock lock(_m);

        if (_destroyed)
            throw ObjectNotExistException(__FILE__, __LINE__);

        c.adapter->remove(id());
        _destroyed = true;
    }

    _parent->removeEntry(_name);
}
```


Copyright 2016, ZeroC, Inc.

Object Life Cycle Concurrency Issues in C++

The preceding implementation is provably deadlock free. All member functions hold only one lock at a time, so they cannot deadlock with each other or themselves. While the locks are held, the functions do not call other member functions that acquire locks, so any potential deadlock can only arise by concurrent calls to another mutating function, either on the same node or on different nodes. For concurrent calls on the same node, deadlock is impossible because such calls are strictly serialized on the mutex `_m`; for concurrent calls to `destroy` on different nodes, each node locks its respective mutex `_m`, releases `_m` again, and then acquires and releases a lock on its parent (by calling `removeEntry`), also making deadlock impossible.

See Also

- Example of a File System Server in C++
- Life Cycle and Normal Operations
- Acquiring Locks without Deadlocks
- Object Identity and Uniqueness
- The Active Servant Map

Copyright 2016, ZeroC, Inc.

Implementing Object Life Cycle in Java

The implementation of our life cycle design has the following characteristics:

- It uses UUIDs as the object identities for nodes to avoid object reincarnation problems.
- When `destroy` is called on a node, the node needs to destroy itself and inform its parent directory that it has been destroyed (because the parent directory is the node's factory and also acts as a collection manager for child nodes).

Note that, in contrast to the initial version, the entire implementation resides in a `FilesystemI` package instead of being part of the `Filesystem` package. Doing this is not essential, but is a little cleaner because it keeps the implementation in a package that is separate from the Slice-generated package.

On this page:

- Object Life Cycle Changes for the NodeI Class in Java
- Object Life Cycle Changes for the DirectoryI Class in Java
- Object Life Cycle Changes for the FileI Class in Java
- Object Life Cycle Concurrency Issues in Java

Object Life Cycle Changes for the `NodeI` Class in Java

Our `DirectoryI` and `FileI` servants derive from a common `NodeI` base interface. This interface is not essential, but useful because it allows us to treat servants of type `DirectoryI` and `FileI` polymorphically:

Java

```
package FilesystemI;

public interface NodeI
{
    Ice.Identity id();
}
```

The only method is the `id` method, which returns the identity of the corresponding node.

Object Life Cycle Changes for the `DirectoryI` Class in Java

As in the initial version, the `DirectoryI` class derives from the generated base class `_DirectoryDisp`. In addition, the class implements the `NodeI` interface. `DirectoryI` must implement each of the Slice operations, leading to the following outline:

Copyright 2016, ZeroC, Inc.

Java

```
package FilesystemI;

import Ice.*;
import Filesystem.*;

public class DirectoryI extends _DirectoryDisp implements NodeI
{
    public Identity
    id();

    public synchronized String
    name(Current c);

    public synchronized NodeDesc[]
    list(Current c);

    public synchronized NodeDesc
    find(String name, Current c) throws NoSuchName;

    public synchronized FilePrx
    createFile(String name, Current c) throws NameInUse;

    public synchronized DirectoryPrx
    createDirectory(String name, Current c) throws NameInUse;

    public void
    destroy(Current c) throws PermissionDenied;

    // ...
}
```

To support the implementation, we also require a number of methods and data members:

Copyright 2016, ZeroC, Inc.

Java

```
package FilesystemI;

import Ice.*;
import Filesystem.*;

public class DirectoryI extends _DirectoryDisp implements NodeI
{
    // ...

    public DirectoryI();
    public DirectoryI(String name, DirectoryI parent);

    public synchronized void
    removeEntry(String name);

    private String _name;       // Immutable
    private DirectoryI _parent; // Immutable
    private Identity _id;       // Immutable
    private boolean _destroyed;
    private java.util.Map<String, NodeI> _contents;
}
```

The `_name` and `_parent` members store the name of this node and a reference to the node's parent directory. (The root directory's `_parent` member is null.) Similarly, the `_id` member stores the identity of this directory. The `_name`, `_parent`, and `_id` members are immutable once they have been initialized by the constructor. The `_destroyed` member prevents a race condition; to interlock access to `_destroyed` (as well as the `_contents` member) we can use synchronized methods (as for the `name` method), or use a `synchronized(this)` block.

The `_contents` map records the contents of a directory: it stores the name of an entry, together with a reference to the child node.

Here are the two constructors for the class:

Java

```
public DirectoryI()
{
    this("/", null);
}

public DirectoryI(String name, DirectoryI parent)
{
    _name = name;
    _parent = parent;
    _id = new Identity();
    _destroyed = false;
    _contents = new java.util.HashMap<String, NodeI>();

    _id.name = parent == null ? "RootDir" : java.util.UUID.randomUUID().toString();
}
```


Copyright 2016, ZeroC, Inc.

The first constructor is a convenience function to create the root directory with the fixed identity "RootDir" and a null parent.

The real constructor initializes the `_name`, `_parent`, `_id`, `_destroyed`, and `_contents` members. Note that nodes other than the root directory use a UUID as the object identity.

The `removeEntry` method is called by the child to remove itself from its parent's `_contents` map:

Java

```
public synchronized void
removeEntry(String name)
{
    _contents.remove(name);
}
```

The implementation of the Slice name operation simply returns the name of the node, but also checks whether the node has been destroyed:

Java

```
public synchronized String
name(Current c)
{
    if (_destroyed)
        throw new ObjectNotExistException();
    return _name;
}
```

Note that this method is synchronized, so the `_destroyed` member cannot be accessed concurrently.

Here is the `destroy` member function for directories:

Copyright 2016, ZeroC, Inc.

Java

```
public void
destroy(Current c) throws PermissionDenied
{
    if (_parent == null)
        throw new PermissionDenied("Cannot destroy root directory");

    synchronized(this) {
        if (_destroyed)
            throw new ObjectNotExistException();

        if (_contents.size() != 0)
            throw new PermissionDenied("Cannot destroy non?empty directory");

        c.adapter.remove(id());
        _destroyed = true;
    }

    _parent.removeEntry(_name);
}
```

The code first prevents destruction of the root directory and then checks whether this directory was destroyed previously. It then acquires the lock and checks that the directory is empty. Finally, `destroy` removes the Active Servant Map (ASM) entry for the destroyed directory and removes itself from its parent's `_contents` map. Note that we call `removeEntry` outside the synchronization to avoid deadlocks.

The `createDirectory` implementation acquires the lock before checking whether the directory already contains a node with the given name (or an invalid empty name). If not, it creates a new servant, adds it to the ASM and the `_contents` map, and returns its proxy:

Java

```
public synchronized DirectoryPrx
createDirectory(String name, Current c) throws NameInUse
{
    if (_destroyed)
        throw new ObjectNotExistException();

    if (name.length() == 0 || _contents.containsKey(name))
        throw new NameInUse(name);

    DirectoryI d = new DirectoryI(name, this);
    ObjectPrx node = c.adapter.add(d, d.id());
    _contents.put(name, d);
    return DirectoryPrxHelper.uncheckedCast(node);
}
```

The `createFile` implementation is identical, except that it creates a file instead of a directory:

Copyright 2016, ZeroC, Inc.

Java

```
public synchronized FilePrx
createFile(String name, Current c) throws NameInUse
{
    if (_destroyed)
        throw new ObjectNotExistException();

    if (name.length() == 0 || _contents.containsKey(name))
        throw new NameInUse(name);

    FileI f = new FileI(name, this);
    ObjectPrx node = c.adapter.add(f, f.id());
    _contents.put(name, f);
    return FilePrxHelper.uncheckedCast(node);
}
```

Here is the implementation of `list`:

Java

```
public synchronized NodeDesc[]
list(Current c)
{
    if(_destroyed)
        throw new ObjectNotExistException();

    NodeDesc[] ret = new NodeDesc[_contents.size()];
    java.util.Iterator<java.util.Map.Entry<String, NodeI> > pos =
        _contents.entrySet().iterator();
    for(int i = 0; i < _contents.size(); ++i) {
        java.util.Map.Entry<String, NodeI> e = pos.next();
        NodeI p = e.getValue();
        ret[i] = new NodeDesc();
        ret[i].name = e.getKey();
        ret[i].type = p instanceof FileI ? NodeType.FileType : NodeType.DirType;
        ret[i].proxy = NodePrxHelper.uncheckedCast(c.adapter.createProxy(p.id()));
    }
    return ret;
}
```

After acquiring the lock, the code iterates over the directory's contents and adds a `NodeDesc` structure for each entry to the returned array.

The `find` operation proceeds along similar lines:

Copyright 2016, ZeroC, Inc.

Java

```
public synchronized NodeDesc
find(String name, Current c) throws NoSuchName
{
    if (_destroyed)
        throw new ObjectNotExistException();

    NodeI p = _contents.get(name);
    if (p == null)
        throw new NoSuchName(name);

    NodeDesc d = new NodeDesc();
    d.name = name;
    d.type = p instanceof FileI ? NodeType.FileType : NodeType.DirType;
    d.proxy = NodePrxHelper.uncheckedCast(c.adapter.createProxy(p.id()));
    return d;
}
```

Object Life Cycle Changes for the `FileI` Class in Java

The `FileI` class is similar to the `DirectoryI` class. The data members store the name, parent, and identity of the file, as well as the `_destroyed` flag and the contents of the file (in the `_lines` member). The constructor initializes these members:

Copyright 2016, ZeroC, Inc.

Java

```
package FilesystemI;

import Ice.*;
import Filesystem.*;
import FilesystemI.*;

public class FileI extends _FileDisp implements NodeI
{
    // ...

    public FileI(String name, DirectoryI parent)
    {
        _name = name;
        _parent = parent;
        _destroyed = false;
        _id = new Identity();
        _id.name = Util.generateUUID();
    }

    private String _name;
    private DirectoryI _parent;
    private boolean _destroyed;
    private Identity _id;
    private String[] _lines;
}
```

The implementation of the remaining member functions of the `FileI` class is trivial, so we present all of them here:

Copyright 2016, ZeroC, Inc.

Java

```
public synchronized String
name(Current c)
{
    if (_destroyed)
        throw new ObjectNotExistException();
    return _name;
}

public Identity
id()
{
    return _id;
}

public synchronized String[]
read(Current c)
{
    if (_destroyed)
        throw new ObjectNotExistException();

    return _lines;
}

public synchronized void
write(String[] text, Current c)
{
    if (_destroyed)
        throw new ObjectNotExistException();

    _lines = (String[])text.clone();
}

public void
destroy(Current c)
{
    synchronized(this) {
        if (_destroyed)
            throw new ObjectNotExistException();

        c.adapter.remove(id());
        _destroyed = true;
    }

    _parent.removeEntry(_name);
}
```

Object Life Cycle Concurrency Issues in Java

The preceding implementation is provably deadlock free. All methods hold only one lock at a time, so they cannot deadlock with each other or themselves. While the locks are held, the methods do not call other methods that acquire locks, so any potential deadlock can only arise

Copyright 2016, ZeroC, Inc.

by concurrent calls to another mutating method, either on the same node or on different nodes. For concurrent calls on the same node, deadlock is impossible because such calls are strictly serialized on the instance; for concurrent calls to `destroy` on different nodes, each node locks itself, releases itself again, and then acquires and releases a lock on its parent (by calling `removeEntry`), also making deadlock impossible.

See Also

- Example of a File System Server in Java
- Life Cycle and Normal Operations
- Object Identity and Uniqueness
- Acquiring Locks without Deadlocks
- The Active Servant Map

Copyright 2016, ZeroC, Inc.

Avoiding Server-Side Garbage

On this page:

- Creating Garbage
- Garbage Collection Solutions
- Simple Mechanisms for Garbage Collection
- Using an Extra Level of Indirection to Collect Garbage
- Server-Side Changes for Garbage Collection
- Client-Side Changes for Garbage Collection

Creating Garbage

We have discussed the implementation of object life cycle, that is, how to correctly provide clients with the means to create and destroy objects. However, throughout this discussion, we have tacitly assumed that clients actually call `destroy` once they no longer need an object. What if this is actually not the case? For example, a client might intend to call `destroy` on an object it no longer needs but crash before it can actually make the call.

To see why this is a realistic (and serious) scenario, consider an on-line retail application. Typically, such an application provides a shopping cart to the client, into which the client can place items. Naturally, the cart and the items it contains will be modelled as server-side objects. The expectation is that, eventually, the client will either finalize or cancel the purchase, at which point the shopping cart and its contents can be destroyed. However, the client may never do that, for example, because it crashes or simply loses interest.

The preceding scenario applies to many different applications and shows up in various disguises. For example, the objects might be session objects that encapsulate security credentials of a client, or might be iterator objects that allow a client to iterate over a collection of values. The key point is that the interactions between client and server are stateful: the server creates state on behalf of a client, holds that state for the duration of several client-server interactions, and expects the client to inform the server when it can clean up that state. If the client never informs the server, whatever resources are associated with that client's state are leaked; these resources are termed *garbage*.

The garbage might be memory, file descriptors, network connections, disk space, or any number of other things. Unless the server takes explicit action, eventually, the garbage will accumulate to the point where the server fails because it has run out of memory, file descriptors, network connections, or disk space.

In the context of Ice, the garbage are servants and their associated resources. In this section, we examine strategies that a server can use avoid drowning in that garbage.

Garbage Collection Solutions

The server is presented with something of a dilemma by garbage objects. The difficulty is not in how to remove the garbage objects (after all, the server knows how to destroy each object), but how to identify whether a particular object is garbage or not. The server knows when a client uses an object (because the server receives an invocation for the object), but the server does not know when an object is no longer of interest to a client (because a dead client is indistinguishable from a slow one).

One approach to dealing with garbage is to avoid creating it in the first place: if all interactions between client and server are stateless, the garbage problem does not arise. Unfortunately, for many applications, implementing this approach is infeasible. The reason is that, in order to turn interactions that are inherently stateful (such as updating a database) into stateless ones, designers are typically forced to keep all the state on the client side, transmit whatever state is required by the server with each remote procedure call, and return the updated state back to the client. In many situations, this simply does not work: for one, the amount of state that needs to be transmitted with each call is often prohibitively large; second, replicating all the state on the client side creates other problems, such as different clients concurrently making conflicting updates to the same data.

The remainder of this section ignores stateless designs. This is not to say that stateless designs are undesirable: where suitable, they can be very effective. However, because many applications simply cannot use them, we focus instead on other ways to deal with garbage.

Mechanisms that identify and reclaim garbage objects are known as garbage collectors. Garbage collectors are well-understood for non-distributed systems. (For example, many programming languages, such as Java and C#, have built-in garbage collectors.)

Non-distributed garbage collectors keep track of all objects that are created, and perform some form of connectivity analysis to determine which objects are still reachable; any objects that are unreachable (that is, objects to which the application code no longer holds any reference) are garbage and are eventually reclaimed.

Unfortunately, for distributed systems, traditional approaches to garbage collection do not work because the cost of performing the connectivity analysis becomes prohibitively large. For example, DCOM provided a distributed garbage collector that turned out to be its Achilles' heel: the collector did not scale to large numbers of objects, particularly across WANs, and several attempts at making it scale failed.

An alternative to distributed garbage collection is to use timeouts to avoid the cost of doing a full connectivity analysis: if an object has not been used for a certain amount of time, the server assumes that it is garbage and reclaims the object. The drawback of this idea is that it is possible for objects to be collected while they are still in use. For example, a customer may have placed a number of items in a shopping

Copyright 2016, ZeroC, Inc.

cart and gone out to lunch, only to find on return that the shopping cart has disappeared in the mean time.

Yet another alternative is to use the evictor pattern: the server puts a cap on the total number of objects it is willing to create on behalf of clients and, once the cap is reached, destroys the least-recently used object in order to make room for a new one. This puts an upper limit on the resources used by the server and eventually gets rid of all unwanted objects. But the drawback is the same as with timeouts: just because an object has not been used for a while does not necessarily mean that it truly is garbage.

Neither timeouts nor evictors are true garbage collectors because they can collect objects that are not really garbage, but they do have the advantage that they reap objects even if the client is alive, but forgets to call `destroy` on some of these objects.

Simple Mechanisms for Garbage Collection

Traditional garbage collection fails in the distributed case for a number of reasons:

- Garbage collectors require connectivity analysis, which is prohibitively expensive. Furthermore, for distributed object systems that permit proxies to be externalized as strings, such as Ice, connectivity analysis is impossible because proxies can exist and travel by means that are invisible to the run time. For example, proxies can exist as records in a database and can travel as strings inside e-mail messages.
- Garbage collectors consider all objects in existence but, for the vast majority of applications, only a small subset of all objects actually ever needs collecting. The work spent in examining objects that can never become garbage is wasted.
- Garbage collectors examine connectivity at the granularity of a single object. However, for many distributed applications, objects are used in groups and, if one object in a group is garbage, all objects in the group are garbage. It would be useful to take advantage of this knowledge, but a garbage collector cannot do this because that knowledge is specific to each application.

In the remainder of this section, we examine a simple mechanism that allows you to get rid of garbage objects cheaply and effectively. The approach has the following characteristics:

- Only those objects that potentially can become garbage are considered for collection.
- Granularity of collection is under control of the application: you can have objects collected as groups of arbitrary size, down to a single object.
- Objects are guaranteed not to be collected prematurely.
- Objects are guaranteed to be collected if the client crashes or suffers loss of connectivity.
- The mechanism is simple to implement and has low run-time overhead.

It is equally important to be aware of the limitations of the approach:

- The approach collects objects if a client crashes, but offers no protection against clients that are still running, but have neglected to destroy objects that they no longer need. In other words, the server is protected against client-side hardware failure and catastrophic client crashes, but it is not protected against faulty programming logic of clients.
- The approach is not transparent at the interface level: it requires changes (albeit minor ones) to the interface definitions for an application.
- The approach requires the client to periodically call the server, thus consuming network resources even if the client is otherwise idle.

Despite the limitations, this approach to garbage collection is applicable to a wide variety of applications and meets the most pragmatic need: how to clean up in case something goes badly wrong (rather than how to clean up in case the client misbehaves).

Using an Extra Level of Indirection to Collect Garbage

Object factories are typically singleton objects that create objects on behalf of various clients. It is important for our garbage collector to know which client created what objects, so the collector can reap the objects created by a specific client if that client crashes. We can easily deal with this requirement by adding the proverbial level of indirection: instead of making a factory a singleton object, we provide a singleton object that creates factories. Clients first create a factory and then create all other objects they need using that factory:

Copyright 2016, ZeroC, Inc.

Slice

```
interface Item { /* ...*/ };

interface Cart
{
    Item* create(/* ... */);
    idempotent string getName();
    void destroy();
    idempotent void refresh();
};

interface CartFactory // Singleton
{
    Cart* create(string name);
};
```

Clients obtain a proxy to the `CartFactory` singleton and call `create` to create a `Cart` object. In turn, the `Cart` object provides a create operation to place new objects of type `Item` into the cart. Note that the shopping cart name allows a client to distinguish different carts — the `name` parameter is *not* used to identify clients or to provide a unique identifier for carts. The `getName` operation on the `Cart` object returns the name that was used by the client to create it.

Each `Cart` object remembers which items it created. Because each client uses its own shopping cart, the server knows which objects were created by what client. In normal operation, a client first creates a shopping cart, and then uses the cart to create the items in the cart. Once the client has finished its job, it calls `destroy` on the cart. The implementation of `destroy` destroys both the cart and its items to reclaim resources.

To deal with crashed clients, the server needs to know when a cart is no longer in use. This is the purpose of the `refresh` operation: clients are expected to periodically call `refresh` on their cart objects. For example, the server might decide that, if a client's cart has not been refreshed for more than ten minutes, the cart is no longer in use and reclaim it. As long as the client calls `refresh` at least once every ten minutes, the cart (and the items it contains) remain alive; if more than ten minutes elapse, the server simply calls `destroy` on the cart. Of course, there is no need to hard-wire the timeout value — you can make it part of the application's configuration. However, to keep the implementation simple, it is useful to have the same timeout value for all carts, or to at least restrict the timeouts for different carts to a small number of fixed choices — this considerably simplifies the implementation in both client and server.

Server-Side Changes for Garbage Collection

The implementation of the interfaces on the server side almost suggests itself:

- Whenever `refresh` is called on a cart, the cart records the time at which the call was made.
- The server runs a reaper thread that wakes up once every ten minutes. The reaper thread examines the timestamp of all carts and, if it finds a cart last time-stamped more than ten minutes ago, it calls `destroy` on that cart.
- Each cart remembers the items it contains and destroys them as part of its `destroy` implementation.

Here then is the reaper thread in outline. (Note that we have simplified the code to show the essentials. For example, we have omitted the code that is needed to make the reaper thread terminate cleanly when the server shuts down. See the code in `demo/Ice/session` for more detail.)

Copyright 2016, ZeroC, Inc.

C++

```
class ReapThread : public IceUtil::Thread,
                   public IceUtil::Monitor<IceUtil::Mutex>
{
public:
    ReapThread();
    virtual void run();
    void add(const CartPrx&, const CartIPtr&);

private:
    const IceUtil::Time _timeout;
    struct CartProxyPair
    {
        CartProxyPair(const CartPrx& p, const CartIPtr& c) :
              proxy(p), cart(c) { }
        const CartPrx proxy;
        const CartIPtr cart;
    };
    std::list<CartProxyPair> _carts;
};

typedef IceUtil::Handle<ReapThread> ReapThreadPtr;
```

Note that the reaper thread maintains a list of pairs. Each pair stores the proxy of a cart and its servant pointer. We need both the proxy and the pointer because we need to invoke methods on both the Slice interface and the implementation interface of the cart. Whenever a client creates a new cart, the server calls the `add` method on the reaper thread, passing it the new cart:

C++

```
void ReapThread::add(const CartPrx& proxy, const CartIPtr& cart)
{
    Lock sync(*this);
    _carts.push_back(CartProxyPair(proxy, cart));
}
```

The `run` method of the reaper thread is a loop that sleeps for ten minutes and calls `destroy` on any session that has not been refreshed within the preceding ten minutes:

Copyright 2016, ZeroC, Inc.

C++

```
void ReapThread::run()
{
    Lock sync(*this);
    while (true) {
        timedWait(_timeout);
        list<CartProxyPair>::iterator p = _carts.begin();
        while (p != _carts.end()) {
            try {
                //
                // Cart destruction may take some time.
                // Therefore the current time is computed
                // for each iteration.
                //
                if ((IceUtil::Time::now() - p->cart->timestamp()) > _timeout) {
                    p->proxy->destroy();
                    p = _carts.erase(p);
                } else {
                    ++p;
                }
            } catch (const Ice::ObjectNotExistException&) {
                p = _carts.erase(p);
            }
        }
    }
}
```

Note that the reaper thread catches `ObjectNotExistException` from the call to `destroy`, and removes the cart from its list in that case. This is necessary because it is possible for a client to call `destroy` explicitly, so a cart may be destroyed already by the time the reaper thread examines it.

The `CartFactory` implementation is trivial:

C++

```
class CartFactoryI : CartFactory
{
public:
    CartFactoryI(const ReapThreadPtr&);
    virtual CartPrx create(const std::string&, const Ice::Current&);

private:
    ReapThreadPtr _reaper;
};
```

The constructor is passed the instantiated reaper thread and remembers that thread in the `_reaper` member.

The `create` method adds each new cart to the reaper thread's list of carts:

Copyright 2016, ZeroC, Inc.

C++

```
CartPrx CartFactoryI::create(const string& name, const Ice::Current& c)
{
    CartIPtr cart = new CartI(name);
    CartPrx proxy = CartPrx::uncheckedCast(
                        c.adapter->addWithUUID(cart));
                        _reaper->add(proxy, cart);
    return proxy;
}
```

Note that each cart internally has a unique ID that is unrelated to its name — the name exists purely as a convenience for the application.

The server's `main` function starts the reaper thread and instantiates the cart factory:

C++

```
ReapThreadPtr reaper = new ReapThread();
CartFactory factory = new CartFactoryI(reaper);
reaper->start();
adapter->add(factory, Ice::stringToIdentity(CartFactory));
adapter->activate();
```

This completes the implementation on the server side. Note that there is very little code here, and that much of this code is essentially the same for each application. For example, we could easily turn the `ReapThread` class into a template class to permit the same code to be used for something other than shopping carts.

Client-Side Changes for Garbage Collection

On the client side, the application code does what it would do with an ordinary factory, except for the extra level of indirection: the client first creates a cart, and then uses the cart as its factory.

As long as the client-side calls `refresh` at least once every ten minutes, the cart remains alive and, with it, all items the client created in that cart. Once the client misses a `refresh` call, the reaper thread in the server cleans up the cart and its items.

To keep the cart alive, you could sprinkle your application code with calls to `refresh` in the hope that at least one of these calls is made at least every ten minutes. However, that is not only error-prone, but also fails if the client blocks for some time. A much better approach is to run a thread in the client that automatically calls `refresh`. That way, the calls are guaranteed to happen even if the client's main thread blocks for some time, and the application code does not get polluted with `refresh` calls. Again, we show a simplified version of the refresh thread here that does not deal with issues such as clean shutdown and a few other irrelevant details:

Copyright 2016, ZeroC, Inc.

C++

```
class CartRefreshThread : public IceUtil::Thread,
                          public IceUtil::Monitor<IceUtil::Mutex>
{
public:
    CartRefreshThread(const IceUtil::Time& timeout, const CartPrx& cart) :
        _cart(cart),
        _timeout(timeout) {}

    virtual void run() {
        Lock sync(*this);
        while(true) {
            timedWait(_timeout);
            try {
                _cart->refresh();
            } catch(const Ice::Exception& ex) {
                return;
            }
        }
    }

private:
    const CartPrx _cart;
    const IceUtil::Time _timeout;
};

typedef IceUtil::Handle<CartRefreshThread> CartRefreshThreadPtr;
```

The client's `main` function instantiates the reaper thread after creating a cart. We assume that the client has a proxy to the cart factory in the `factory` variable:

C++

```
CartPrx cart = factory->create(name);
CartRefreshThreadPtr refresh = new CartRefreshThread(IceUtil::Time::seconds(480),
cart);
refresh->start();
```

Note that, to be on the safe side and also allow for some network delays, the client calls refresh every eight minutes; this is to ensure that at least one call to `refresh` arrives at the server within each ten-minute interval.

See Also

- Object Life Cycle

Copyright 2016, ZeroC, Inc.

Platform-Specific Features

This section describes Ice features that are available only on a specific platform, such as Microsoft Windows.

Topics

- Windows Services
- Windows Store Applications

Copyright 2016, ZeroC, Inc.

Windows Services

A Windows service is a program that runs in the background and typically does not require user intervention. Similar to a daemon on Unix platforms, a Windows service is usually launched automatically when the operating system starts and runs until the system shuts down.

Ice includes the `Service` class that simplifies the task of writing an Ice-based Windows service in C++. Writing the service is only the first step, however, as it is also critically important that the service be installed and configured correctly for successful operation. Service installation and configuration is outside the scope of the `Service` class because these tasks are generally not performed by the service itself but rather as part of a larger administrative effort to deploy an application. Furthermore, there are security implications to consider when a service is able to install itself: such a service typically needs administrative rights to perform the installation, but does not necessarily need those rights while it is running. A better strategy is to grant administrative rights to a separate installer program and not to the service itself.

Topics

- Installing a Windows Service
- Using the Ice Service Installer
- Manually Installing a Service as a Windows Service
- Troubleshooting Windows Services

Copyright 2016, ZeroC, Inc.

Installing a Windows Service

The installation of a Windows service varies in complexity with the needs of the application, but usually involves the following activities:

- Selecting the user account in which the service will run.
- Registering the service and establishing its activation mode and dependencies.
- Creating one or more file system directories to contain executables, libraries, and supporting files or databases.
- Configuring those directories with appropriate permissions so that they are accessible to the user account selected for the service.
- Creating keys in the Windows registry.
- Configuring the Windows Event Log so that the service can report status and error messages.

There are many ways to perform these tasks. For example, an administrator can execute them manually. Another option is to write a script or program tailored to the needs of your application. Finally, you can build an installer using a developer tool such as InstallShield.

Selecting a User Account for the Service

Before installing a service, you should give careful consideration to the user account that will run the service. Unless your service has special requirements, we recommend that you use the built-in account that Windows provides specifically for this purpose, `Local Service`.

See Also

- Manually Installing a Service as a Windows Service

Copyright 2016, ZeroC, Inc.

Using the Ice Service Installer

Ice provides the command-line tool `iceserviceinstall` to assist you in installing and uninstalling the following Ice services as Windows services:

- IceGrid registry
- IceGrid node
- Glacier2 router

Ice includes other programs that can also be run as Windows services, such as the IceBox and IcePatch2 servers. Typically it is not necessary to install these programs as Windows services because they can be launched by an IceGrid node service. However, if you wish to run an IceBox or IcePatch2 server as a Windows service without the use of IceGrid, you must manually install the service.

Here we describe how to use the Ice service installer and discuss its actions and prerequisites.

On this page:

- iceserviceinstall Command Line Options
- Security Considerations for Ice Services
- iceserviceinstall Configuration File
  - Sample Configuration Files
- iceserviceinstall Properties
- Service Installation Process
- Uninstalling a Windows Service

`iceserviceinstall` Command Line Options

`iceserviceinstall` supports the following options and arguments:

```
iceserviceinstall [options] service config-file [property ...]

Options:
-h, --help           Show this message.
-n, --nopause        Do not call pause after displaying a message.
-v, --version        Display the Ice version.
-u, --uninstall      Uninstall the Windows service.
```

The *`service`* and *`config-file`* arguments are required during installation and uninstallation.

The *`service`* argument selects the type of service you are installing; use one of the following values:

- `icegridregistry`
- `icegridnode`
- `glacier2router`

Note that the Ice service installer currently does not support the installation of an IceGrid node with a collocated registry, therefore you must install the registry and node separately.

The *`config-file`* argument specifies the name of an Ice configuration file.

When installing a service, properties can be defined on the command line using the *`--name=value`* syntax, or they can be defined in the configuration file. The supported properties are described below.

Security Considerations for Ice Services

None of the Ice services require privileges beyond a normal user account. In the case of the IceGrid node service in particular, we do not recommend running it in a user account with elevated privileges because the service is responsible for launching server executables, and those servers would inherit the node's access rights.

`iceserviceinstall` Configuration File

The Ice service installer requires that you specify the path name of the Ice configuration file for the service being installed or uninstalled. The

Copyright 2016, ZeroC, Inc.

tool needs this path name for several reasons:

- During installation, it verifies that the configuration file has sufficient access rights.
- It configures a newly-installed service to load the configuration file using its absolute path name, therefore you must decide in advance where the file will be located.
- It reads the configuration file and examines certain service-specific properties. For example, prior to installing an IceGrid registry service, the tool verifies that the directory specified by the property `IceGrid.Registry.Data` has sufficient access rights.
- The tool supports its own configuration parameters that may also be defined as properties in this file.

You may still modify a service's configuration file after installation, but you should uninstall and reinstall the service if you change any of the properties that influence the service installer's actions. The table below describes the service properties that affect the installer:

Property	Service	Description
`IceGrid.InstanceName`	IceGrid Registry	Value appears in the service name; also included in the default display name if one is not defined.
`IceGrid.Node.Data`	IceGrid Node	Directory is created if necessary; access rights are verified.
`IceGrid.Node.Name`	IceGrid Node	Value appears in the service name; also included in the default display name if one is not defined.
`IceGrid.Registry.Data`	IceGrid Registry	Directory is created if necessary; access rights are verified.
`Ice.Default.Locator`	IceGrid Node, Glacier2 Router	The IceGrid instance name is derived from the identity in this proxy.
`Ice.EventLog.Source`	All	Specifies the name of an event log source for the service.

The steps performed by the tool during an installation are described in detail below.

Sample Configuration Files

Ice includes sample configuration files for the IceGrid and Glacier2 services in the `config` subdirectory of your Ice installation. We recommend that you review the comments and settings in these files to familiarize yourself with a typical configuration of each service.

You can modify a configuration file to suit your needs or copy one to use as a starting point for your own configuration.

`iceserviceinstall` Properties

The Ice service installer uses a set of optional properties that customize the installation process. These properties can be defined in the service's configuration file as discussed above, or they can be defined on the command line using the familiar --*name*=*value* syntax:

```
iceserviceinstall --DependOnRegistry=1 ...
```

The installer's properties are listed below:

- `AutoStart=`*num*
  If num is a value greater than zero, the service is configured to start automatically at system start up. Set this property to zero to configure the service to start on demand instead. If not specified, the default value is `1`.
- `DependOnRegistry=`*num*
  If num is a value greater than zero, the service is configured to depend on the IceGrid registry, meaning Windows will start the registry prior to starting this service. Enabling this feature also requires that the property `Ice.Default.Locator` be defined in *config-file*. If not specified, the default value is zero.
- `Description=`*value*
  A brief description of the service. If not specified, a general description is used.
- `DisplayName=`*name*
  The friendly name that identifies the service to the user. If not specified, `iceserviceinstall` composes a default display name.
- `EventLog=`*name*
  The name of the event log used by the service. If not specified, the default value is `Application`.

Copyright 2016, ZeroC, Inc.

- `ImagePath=`*`path`*
  The path name of the service executable. If not specified, `iceserviceinstall` assumes the service executable resides in the same directory as itself and fails if the executable is not found.
- `ObjectName=`*`name`*
  Specifies the account used to run the service. If not specified, the default value is `NT Authority\LocalService`.
- `Password=`*`value`*
  The password required by the account specified in `ObjectName`.

Service Installation Process

The Ice service installer performs a number of steps to install a service. As discussed earlier, you must specify the path name of the service's configuration file because the service installer uses certain properties during the installation process. The actions taken by the service installer are described below:

- Obtain the service's *instance name* from the configuration file. The instance name is specified by the property `IceGrid.InstanceName` or `Glacier2.InstanceName`. If an instance name is not specified, the default value is `IceGrid` or `Glacier2`, respectively. If the service being installed depends on the IceGrid registry, the IceGrid instance name is derived from the value of the `Ice.Default.Locator` property.
- For an IceGrid node, obtain the node's name from the property `IceGrid.Node.Name`. This property must be defined when installing a node.
- Compose the service name from the service type, instance name, and node name (for an IceGrid node). For example, the default service name for an IceGrid registry is `icegridregistry.IceGrid`. Note that the service name is not the same as the display name.
- Resolve the user account specified by `ObjectName`.
- For an IceGrid registry, create the data directory specified by the property `IceGrid.Registry.Data` and ensure that the user account specified by `ObjectName` has read/write access to the directory.
- For an IceGrid node, create the data directory specified by the property `IceGrid.Node.Data` and ensure that the user account specified by `ObjectName` has read/write access to the directory.
- For an IceGrid node, ensure that the user account specified by `ObjectName` has read access to the following registry key:
  `HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib`
  This key allows the node to access CPU utilization statistics.
- Ensure that the user account specified by `ObjectName` has read/write access to the configuration file.
- Create a new Windows event log by adding the registry key specified by `EventLog`.
- Add an event log source under `EventLog` for the source name specified by `Ice.EventLog.Source`. If this property is not defined, the service name is used as the source name.
- Install the service, including command line arguments that specify the service name (`--service` *`name`*) and the absolute path name of the configuration file (`--Ice.Config=`*`config-file`*).

The Ice service installer currently does *not* perform these tasks:

- Modify access rights for the service's executable or its DLL dependencies
- Verify that the user account specified by `ObjectName` has the right to "Log on as a service"

Uninstalling a Windows Service

When uninstalling an existing service, the Ice service installer first ensures that the service is stopped, then proceeds to remove the service. The service's event log source is removed and, if the service is not using the `Application` log, the event log registry key is also removed.

See Also

- icegridregistry
- icegridnode
- Getting Started with Glacier2
- IceBox
- IcePatch2
- Installing a Windows Service
- Manually Installing a Service as a Windows Service
- Troubleshooting Windows Services
- IceGrid.*
- Glacier2.*

Copyright 2016, ZeroC, Inc.

Manually Installing a Service as a Windows Service

This page describes how to manually install and configure an Ice service as a Windows Service using the IcePatch2 service as a case study. For the purposes of this discussion, we assume that Ice is installed in the directory `C:\Ice`. We also assume that you have administrative access to your system, which is required by many of the installation steps discussed below.

On this page:

- Selecting a User Account for the IcePatch2 Service
- Preparing a Directory for the IcePatch2 Service
  - Selecting a Directory for the IcePatch2 Service
  - Creating the Directory for the IcePatch2 Service
  - Populating the Directory for the IcePatch2 Service
- Configuration File for the IcePatch2 Service
- Creating the IcePatch2 Service
- Creating the Event Log for the IcePatch2 Service
  - Using the Application Log for the IcePatch2 Service
  - Using a Custom Log for the IcePatch2 Service
  - Registry Caching for the IcePatch2 Service
- Starting the IcePatch2 Service
- Testing the IcePatch2 Service

Selecting a User Account for the IcePatch2 Service

The IcePatch2 service can run in a regular user account, therefore we will follow our own recommendation and use the Local Service account.

Preparing a Directory for the IcePatch2 Service

The service needs a directory in which to store the files that it distributes to clients. A common mistake is assuming that a service will be able to access a file or directory that you created using your current account, which is likely to cause the service to fail in a way that is difficult to diagnose. To prevent such failures, we will ensure that the directory has the necessary permissions for the service to access it while running in the `Local Service` account.

Selecting a Directory for the IcePatch2 Service

The directory tree for our IcePatch2 service is shown below:

```
C:\Windows\
  ServiceProfiles\
    LocalService\
      AppData\
        Local\
          ZeroC\
            icepatch2\
              data\
```

Creating the Directory for the IcePatch2 Service

Since `Local Service` is a built-in account, its user directory should already exist and have the proper access rights.

> Opening the `LocalService` directory in Windows Explorer may require administrative permission. Once opened, the `AppData` subdirectory may not be visible without modifying the folder options to show hidden files and directories.

You can verify its access rights in a command window using the `icacls` utility:

Copyright 2016, ZeroC, Inc.

```
> cd \Windows\ServiceProfiles
> icacls LocalService
LocalService NT AUTHORITY\SYSTEM:(OI)(CI)(F)
             BUILTIN\Administrators:(OI)(CI)(F)
             NT AUTHORITY\LOCAL SERVICE:(OI)(CI)(F)
```

The `AppData` subdirectory should inherit these same rights. Now we can create the remaining subdirectories:

```
> cd LocalService\AppData
> mkdir Local\ZeroC\icepatch2\data
```

If you want to further restrict access to files or subdirectories, you can modify them as necessary using the `icacls` utility. Note however that certain actions may cause a file to revert back to the access rights of its enclosing directory. For example, modifying a file using a text editor is often the equivalent of erasing the file and recreating it, which discards any access rights you may have previously set for the file.

You can also manage privilege settings interactively using Windows Explorer. For example, right click on the `LocalService` directory, select Properties, and select the Security tab. Click the Advanced button to modify the settings.

Populating the Directory for the IcePatch2 Service

Now you can copy the files that will be distributed to clients into the `data` subdirectory. The new files should inherit the access rights of their enclosing directory. For the sake of discussion, let's copy some Slice files from the Ice distribution into the data directory:

```
> cd Local\ZeroC\icepatch2\data
> copy \Ice\slice\Ice\*.ice
```

Next we need to run `icepatch2calc` to prepare the directory for use by the IcePatch2 service:

```
> cd Local\ZeroC\icepatch2\data
> icepatch2calc .
```

Configuration File for the IcePatch2 Service

IcePatch2 requires a minimal set of configuration properties. We could specify them on the service's command line, but if we later want to modify those properties we would have to reinstall the service. Defining the properties in a file simplifies the task of modifying the service's configuration.

Our IcePatch2 configuration is quite simple:

```
IcePatch2.Directory=C:\Windows\ServiceProfiles\LocalService\AppData\Local\ZeroC\ice
patch2\data
IcePatch2.Endpoints=tcp -p 10000
```

The `IcePatch2.Directory` property specifies the location of the server's data directory, which we created in the previous section.

We will save our configuration properties into the following file:

```
C:\Ice\config\icepatch2.cfg
```


Copyright 2016, ZeroC, Inc.

We must also ensure that the service has permission to access its configuration file. The Ice run time never modifies a configuration file, therefore read access is sufficient. The configuration file likely already has the necessary access rights, which we can verify using the `icacls` utility that we described earlier:

```
> icacls C:\Ice\config\icepatch2.cfg
```

Creating the IcePatch2 Service

We will use Microsoft's Service Control (`sc`) utility in a command window to create the service.

See http://support.microsoft.com/kb/251192 for more information about the SC utility.

Our first `sc` command does the majority of the work (the command is formatted for readability but must be typed on a single line):

```
> sc create icepatch2 binPath= "C:\Ice\bin\icepatch2server.exe
  --Ice.Config=C:\Ice\config\icepatch2.cfg --service icepatch2"
  DisplayName= "IcePatch2 Server" start= auto
  obj= "NT Authority\LocalService" password= ""
```

There are several important aspects of this command:

- The service name is `icepatch2`. You can use whatever name you like, as long as it does not conflict with an existing service. Note however that this name is used in other contexts, such as in the `--service` option discussed below, therefore you must use it consistently.
- Following the service are several options. Note that all of the option names end with an equals sign and are separated from their arguments with at least one space.
- The `binPath=` option is required. We supply the full path name of the IcePatch2 server executable, as well as command-line arguments that define the location of the configuration file and the name of the service, all enclosed in quotes.
- The `DisplayName=` option sets a friendly name for the service.
- The `start=` option configures the start up behavior for the service. We used the argument `auto` to indicate the service should be started automatically when Windows boots.
- The `obj=` option selects the user account in which this service runs. As we explained, the `Local Service` account is appropriate for most services.
- The `password=` option supplies the password associated with the user account indicated by `obj=`. The `Local Service` account has an empty password.

The `sc` utility should report success if it was able to create the service as specified. You can verify that the new service was created with this command:

```
> sc qc icepatch2
```

Alternatively, you can start the Services administrative control panel and inspect the properties of the IcePatch2 service.

If you start the control panel, you will notice that the entry for IcePatch2 does not have a description. To add a description for the service, use the following command:

```
> sc description icepatch2 "IcePatch2 file server"
```

After refreshing the list of services, you should see the new description.

Creating the Event Log for the IcePatch2 Service

By default, programs such as the IcePatch2 service that utilize the `Service` class log messages to the `Application` event log. Below we describe how to prepare the Windows registry for the service's default behavior, and we also show how to use a custom event log instead.

Copyright 2016, ZeroC, Inc.

We make use of Microsoft's Registry (`reg`) utility to modify the registry, although you could also use the interactive `regedit` tool. As always, caution is recommended whenever you modify the registry.

Using the Application Log for the IcePatch2 Service

We must configure an event log source for events to display properly. The first step is to create a registry key with the name of the source. Since the `Service` class uses the service name as the source name by default, we add the key `icepatch2` as shown below:

```
> reg add HKLM\SYSTEM\CurrentControlSet\Services\EventLog\Application\icepatch2
```

Inside this key we must add a value specifies the location of the Ice run time DLL:

```
> reg add HKLM\SYSTEM\CurrentControlSet\Services\EventLog\Application\icepatch2 /v
EventMessageFile /t REG_EXPAND_SZ /d C:\Ice\bin\ice36.dll
```

We will also add a value indicating the types of events that the source supports:

```
> reg add HKLM\SYSTEM\CurrentControlSet\Services\EventLog\Application\icepatch2 /v
TypesSupported /t REG_DWORD /d 7
```

The value 7 corresponds to the combination of the following event types:

- `EVENTLOG_ERROR_TYPE`
- `EVENTLOG_WARNING_TYPE`
- `EVENTLOG_INFORMATION_TYPE`

You can verify that the registry values have been defined correctly using the following command:

```
> reg query HKLM\SYSTEM\CurrentControlSet\Services\EventLog\Application\icepatch2
```

Our configuration of the event log is now complete.

Changing the Source Name for the IcePatch2 Service

Using the configuration described in the previous section, events logged by the IcePatch2 service are recorded in the event log using the source name `icepatch2`. If you prefer to use a source name that differs from the service name, you can replace `icepatch2` in the registry commands with the name of your choosing, but you must also add a matching definition for the property `Ice.EventLog.Source` to the service's configuration file.

For example, to use the source name `Ice File Patching Service`, you would add the registry key as shown below:

```
> reg add "HKLM\SYSTEM\CurrentControlSet\Services\EventLog\Application\Ice File
Patching Service"
```

The commands to add the `EventMessageFile` and `TypesSupported` values must be modified in a similar fashion. Finally, add the following configuration property to `icepatch2.cfg`:

```
Ice.EventLog.Source=Ice File Patching Service
```

Using a Custom Log for the IcePatch2 Service

Copyright 2016, ZeroC, Inc.

You may decide that you want your services to record messages into an application-specific log instead of the `Application` log that is shared by other unrelated services. As an example, let us create a log named `MyApp`:

```
> reg add "HKLM\SYSTEM\CurrentControlSet\Services\EventLog\MyApp"
```

Next we add a subkey for the IcePatch2 service. As described in the previous section, we will use a friendlier source name:

```
> reg add "HKLM\SYSTEM\CurrentControlSet\Services\EventLog\MyApp\Ice File Patching
Service"
```

Now we can define values for `EventMessageFile` and `TypesSupported`:

```
> reg add "HKLM\SYSTEM\CurrentControlSet\Services\EventLog\MyApp\Ice File Patching
Service" /v EventMessageFile /t REG_EXPAND_SZ /d C:\Ice\bin\ice36.dll

> reg add "HKLM\SYSTEM\CurrentControlSet\Services\EventLog\MyApp\Ice File Patching
Service" /v TypesSupported /t REG_DWORD /d 7
```

Finally, we define `Ice.EventLog.Source` in the IcePatch2 service's configuration file:

```
Ice.EventLog.Source=Ice File Patching Service
```

Note that you must restart the Event Viewer control panel after adding the `MyApp` registry key in order to see the new log.

Registry Caching for the IcePatch2 Service

The first time a service logs an event, Windows' Event Log service caches the registry entries associated with the service's source. If you wish to modify a service's event log configuration, such as changing the service to use a custom log instead of the `Application` log, you should perform the following steps:

1. Stop the service.
2. Remove the unwanted event log registry key.
3. Add the new event log registry key(s).
4. Restart the system (or at least the Event Log service).
5. Start the service and verify that the log entries appear in the intended log.

After following these steps, open a log entry and ensure that it displays properly. If it does not, for example if the event properties indicate that the description of an event cannot be found, the problem is likely due to a misconfigured event source. Verify that the value of `EventMessageFile` refers to the correct location of the Ice run time DLL, and that the service is defining `Ice.EventLog.Source` in its configuration file (if necessary).

Starting the IcePatch2 Service

We are at last ready to start the service. In a command window, you can use the `sc` utility:

```
> sc start icepatch2
```

The program usually responds with status information indicating that the start request is pending. You can query the service's status to verify that it started successfully:

Copyright 2016, ZeroC, Inc.

```
> sc query icepatch2
```

The service should now be in the running state. If it is not in this state, open the Event Viewer control panel and inspect the relevant log for more information that should help you to locate the problem. Even if the service started successfully, you may still want to use the Event Viewer to confirm that the service is using the log you expected.

Testing the IcePatch2 Service

Ice includes a graphical IcePatch2 client in the `demo/IcePatch2/MFC` directory of the Ice distribution. Once you have built the client, you can use it to test that the service is working properly.

See Also

- IcePatch2
- `Ice::Service` Class
- Installing a Windows Service

Copyright 2016, ZeroC, Inc.

Troubleshooting Windows Services

This page describes how to troubleshoot Windows Services.

On this page:

- Missing Libraries for a Windows Service
- Windows Firewall Interference
- IceGrid Node Performance Monitoring Issues

Missing Libraries for a Windows Service

One failure that commonly occurs when starting a Windows service is caused by missing DLLs, which usually results in an error window stating a particular DLL cannot be found. Fixing this problem can often be a trial-and-error process because the DLL mentioned in the error may depend on other DLLs that are also missing. It is important to understand that a Windows service is launched by the operating system and can be configured to execute as a different user, which means the service's environment (most importantly its `PATH`) may not match yours and therefore extra steps are necessary to ensure that the service can locate its required DLLs.

The command-line utility `dumpbin` can be used to discover the dependencies of an executable or DLL.

The simplest approach is to copy all of the necessary DLLs to the directory containing the service executable. If this solution is undesirable, another option is to modify the system `PATH` to include the directory or directories containing the required DLLs. (Note that modifying the system `PATH` requires restarting the system.) Finally, you can copy the necessary DLLs to `\WINDOWS\system32`, although we do not recommend this approach.

Copying DLLs to `\WINDOWS\system32` often results in subtle problems later when trying to develop using newer versions of the DLLs. Inevitably you will forget about the DLLs in `\WINDOWS\system32` and struggle to determine why your application is misbehaving or failing to start.

Assuming that DLL issues are resolved, a Windows service can fail to start for a number of other reasons, including

- invalid command-line arguments or configuration properties
- inability to access necessary resources such as file systems and databases, because either the resources do not exist or the service does not have sufficient access rights to them
- networking issues, such as attempting to open a port that is already in use, or DNS lookup failures

Failures encountered by the Ice run time prior to initialization of the communicator are reported to the Windows event log if no other logger implementation is defined, so that should be the first place you look. Typically you will find an entry in the `System` event log resembling the following message:

```
The IcePatch2 service terminated with service-specific error 1.
```

Error code `1` corresponds to `EXIT_FAILURE`, the value used by the `Service` class to indicate a failure during startup. Additional diagnostic messages may be available in the `Application` event log. See `Ice::Service` Logging Considerations for more information on configuring a logger for a Windows service.

As we mentioned earlier, insufficient access rights can also prevent a Windows service from starting successfully. By default, a Windows service is configured to run under a local system account, in which case the service may not be able to access resources owned by other users. It may be necessary for you to configure a service to run under a different account, which you can do using the Services control panel. You should also review the access rights of files and directories required by the service.

Windows Firewall Interference

Your choice of user account determines whether you receive any notification when the Windows Firewall blocks the ports that are used by your service. For example, if you use `Local Service` as we recommended, you will not see any Windows Security Alert dialog (see this Microsoft article for details).

If you are not prompted to unblock your service, you will need to manually add an exception in Windows Firewall. For example, follow the steps below to unblock the ports of a Glacier2 router service:

1. Open the Windows Firewall Settings panel and navigate to the Exceptions panel.

Copyright 2016, ZeroC, Inc.

2. Select "Add program..."
3. Select "Browse," navigate to the Glacier2 router executable, and click "OK."

Note that adding an exception for a program unblocks all ports on which the program listens. Review the endpoint configurations of your services carefully to ensure that no unnecessary ports are opened.

For services listening on one or a few fixed ports, you could also create port exceptions in your Windows Firewall. Refer to the Windows Firewall documentation for details.

IceGrid Node Performance Monitoring Issues

The IceGrid node uses Windows' `Perflib` facility to obtain statistics about the CPU utilization of its host for load balancing purposes. Occasionally, the IceGrid node may log the following warning message:

```
warning: Unable to lookup the performance counter name
```

This message is an indication that the node does not have sufficient privileges to access a key in the Windows registry:

```
HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib
```

As part of its installation procedure, the `iceserviceinstall` utility modifies the permissions of this registry key to grant read access to the node's designated user account. If you are trying to change the node's user account, we recommend using the `iceserviceinstall` utility to uninstall and reinstall the node. If you wish to modify the permissions of this registry key manually, follow these steps:

1. Start `regedit` and navigate to the `Perflib` key.
2. Right click on `Perflib` and select `Permissions`.
3. If the desired user account is not already present, click `Add` to add the user account. Enter `LOCAL SERVICE` if you wish to run the node in the Local Service account, otherwise enter the name of the user account. Press `OK`.
4. Check the `Read` box in the `Allow` column to grant read access to the registry key and press `OK` to apply the changes.

Another way to grant the node's user account with the necessary access rights is to add it to the `Performance Monitor Users` group.

See Also

- Load Balancing
- Installing a Windows Service
- Manually Installing a Service as a Windows Service

Copyright 2016, ZeroC, Inc.

Windows Store Applications

A Windows Store application that uses Ice must request certain capabilities in its manifest in order to access public and private networks. The image below shows the relevant settings in Visual Studio:

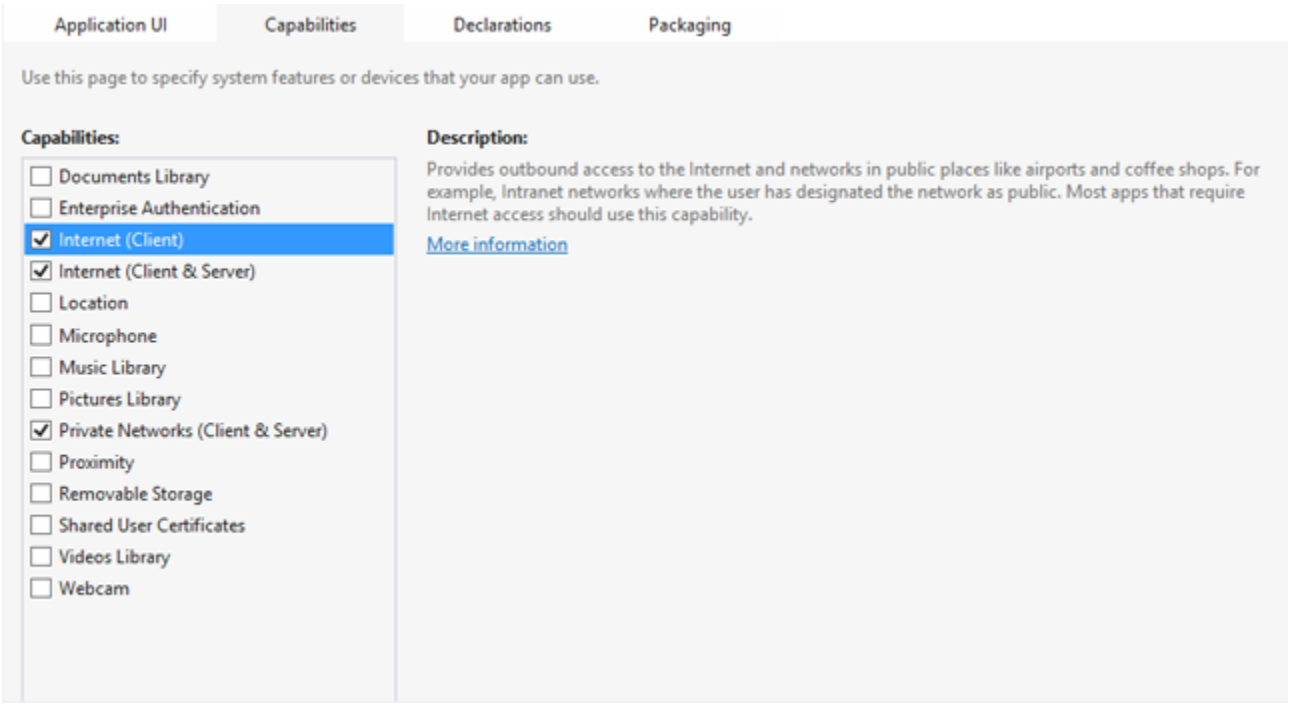

Neglecting to enable the proper capabilities can raise `Ice::SocketException` or, if you are debugging the application in Visual Studio at the time of the exception, it will likely break on `Platform::AccessDeniedException`.

If your client application uses SSL, you can bundle certificates with the application by including them in the application's declarations:

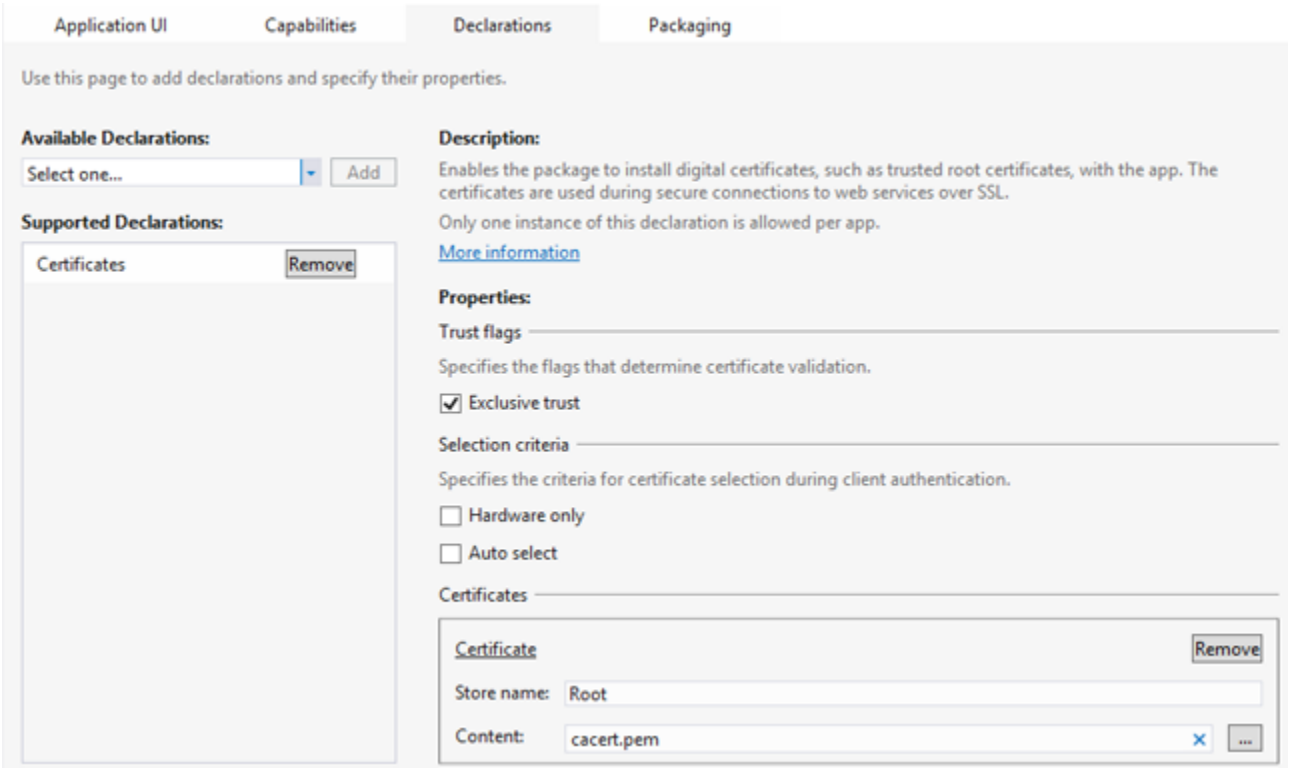

Copyright 2016, ZeroC, Inc.

Example
The UWP sample programs in the Ice distribution use capabilities and declarations.

Copyright 2016, ZeroC, Inc.

Property Reference

This section provides a reference for all properties used by the Ice run time and its services.

Unless stated otherwise in the description of an individual property, its default value is the empty string. If a property takes a numeric value, the empty string is interpreted as zero.

Note that Ice reads properties that control the run time and its services only once on start-up, when you create a communicator. This means that you must set Ice-related properties to their correct values before you create a communicator. If you change the value of an Ice-related property after that point, it is likely that the new setting will simply be ignored.

Topics

- Object Adapter Properties
- Proxy Properties
- Miscellaneous Ice.* Properties
- Freeze.*
- Glacier2.*
- Ice.ACM.*
- Ice.Admin.*
- Ice.Config
- Ice.Default.*
- Ice.InitPlugins
- Ice.IPv4
- Ice.IPv6
- Ice.Override.*
- Ice.Plugin.*
- Ice.PluginLoadOrder
- Ice.PreferIPv6Address
- Ice.TCP.*
- Ice.ThreadPool.*
- Ice.ThreadPriority
- Ice.Trace.*
- Ice.UDP.*
- Ice.Warn.*
- IceBox.*
- IceBoxAdmin
- IceDiscovery.*
- IceGrid.*
- IceGridAdmin.*
- IceLocatorDiscovery.*
- IceMX.Metrics.*
- IcePatch2.*
- IcePatch2Client.*
- IceSSL.*
- IceStorm Properties
- IceStormAdmin.*

Copyright 2016, ZeroC, Inc.

Object Adapter Properties

On this page:

- adapter.ACM.Close
- adapter.ACM.Heartbeat
- adapter.ACM.Timeout
- adapter.AdapterId
- adapter.Endpoints
- adapter.Locator
- adapter.MessageSizeMax
- adapter.ProxyOptions
- adapter.PublishedEndpoints
- adapter.ReplicaGroupId
- adapter.Router
- adapter.ThreadPool.Serialize
- adapter.ThreadPool.Size
- adapter.ThreadPool.SizeMax
- adapter.ThreadPool.SizeWarn
- adapter.ThreadPool.StackSize
- adapter.ThreadPool.ThreadIdleTime
- adapter.ThreadPool.ThreadPriority

adapter.ACM.Close

Synopsis

`adapter.ACM.Close=num`

Description

Overrides the value of `Ice.ACM.Server.Close` for incoming connections to the named object adapter.

Note that ACM can cause incoming oneway requests to be silently discarded.

adapter.ACM.Heartbeat

Synopsis

`adapter.ACM.Heartbeat=num`

Description

Overrides the value of `Ice.ACM.Server.Heartbeat` for incoming connections to the named object adapter.

adapter.ACM.Timeout

Synopsis

`adapter.ACM.Timeout=num`

Description

Overrides the value of `Ice.ACM.Server.Timeout` for incoming connections to the named object adapter.

adapter.AdapterId

Copyright 2016, ZeroC, Inc.

Synopsis

adapter.AdapterId=*id*

Description

Specifies an identifier for the object adapter with the name *adapter*. This identifier must be unique among all object adapters using the same locator instance. If a locator proxy is defined using *adapter*.Locator or Ice.Default.Locator, this object adapter sets its endpoints with the locator registry upon activation.

adapter.Endpoints

Synopsis

adapter.Endpoints=*endpoints*

Description

Sets the endpoints for the object adapter *adapter* to *endpoints*. These endpoints specify the network interfaces on which the object adapter receives requests. Proxies created by the object adapter contain these endpoints, unless the *adapter*.PublishedEndpoints pro perty is also specified.

adapter.Locator

Synopsis

adapter.Locator=*locator*

Description

Specifies a locator for the object adapter with the name *adapter*. The value is a stringified proxy to the Ice locator interface.

As a proxy property, you can configure additional aspects of the proxy using properties.

adapter.MessageSizeMax

Synopsis

adapter.MessageSizeMax=*num*

Description

Overrides the setting of Ice.MessageSizeMax to limit the size of messages that can be received by this object adapter. If not defined, the adapter uses the value of Ice.MessageSizeMax.

adapter.ProxyOptions

Synopsis

adapter.ProxyOptions=*options*

Description

Specifies the proxy options for proxies created by the object adapter. The value is a string representing the proxy options as they would be specified in a stringified proxy. The default value is "-t", that is, proxies created by the object adapter are configured to use twoway invocations by default.

Copyright 2016, ZeroC, Inc.

adapter.PublishedEndpoints

Synopsis

`adapter.PublishedEndpoints=endpoints`

Description

When creating a proxy, the object adapter *adapter* normally includes the endpoints defined by *adapter*`.Endpoints`. If *adapter*`.PublishedEndpoints` is defined, the object adapter publishes these endpoints instead. This is useful in many situations, such as when a server resides behind a port-forwarding firewall, in which case the object adapter's public endpoints must specify the address and port of the firewall. The *adapter*`.ProxyOptions` property also influences the proxies created by an object adapter.

adapter.ReplicaGroupId

Synopsis

`adapter.ReplicaGroupId=id`

Description

Identifies the group of replicated object adapters to which this adapter belongs. The replica group is treated as a virtual object adapter, so that an indirect proxy of the form `identity@`*id* refers to the object adapters in the group. During binding, a client will attempt to establish a connection to an endpoint of one of the participating object adapters, and automatically try others until a connection is successfully established or all attempts have failed. Similarly, an outstanding request will, when permitted, automatically fail over to another object adapter of the replica group upon connection failure. The set of endpoints actually used by the client during binding is determined by the locator's configuration policies.

Defining a value for this property has no effect unless *adapter*`.AdapterId` is also defined. Furthermore, the locator registry may require replica groups to be defined in advance (see `IceGrid.Registry.DynamicRegistration`), otherwise `Ice.NotRegisteredException` is raised upon adapter activation. Regardless of whether an object adapter is replicated, it can always be addressed individually in an indirect proxy if it defines a value for *adapter*`.AdapterId`.

adapter.Router

Synopsis

`adapter.Router=router`

Description

Specifies a router for the object adapter with the name *adapter*. The value is a stringified proxy to the Ice router control interface. Defining a router allows the object adapter to receive callbacks from the router over a bidirectional connection, thereby avoiding the need for the router to establish a connection back to the object adapter.

A router can only be assigned to one object adapter. Specifying the same router for more than one object adapter results in undefined behavior. The default value is no router.

As a proxy property, you can configure additional aspects of the proxy using properties.

adapter.ThreadPool.Serialize

Synopsis

`adapter.ThreadPool.Serialize=num`

Description

If *num* is a value greater than zero, the adapter's thread pool serializes all messages from each connection. It is not necessary to enable this

Copyright 2016, ZeroC, Inc.

feature in a thread pool whose maximum size is one thread. In a multi-threaded pool, enabling serialization allows requests from different connections to be dispatched concurrently while preserving the order of messages on each connection. Note that serialization has a signficant impact on latency and throughput. If not defined, the default value is zero.

adapter.ThreadPool.Size

Synopsis

`adapter.ThreadPool.Size=num`

Description

A communicator creates a default server thread pool that dispatches requests to its object adapters. An object adapter can also be configured with its own thread pool. This is useful in avoiding deadlocks due to thread starvation by ensuring that a minimum number of threads is available for dispatching requests to certain Ice objects.

num is the initial number of threads in the thread pool. The default value is zero, meaning that an object adapter by default uses the communicator's server thread pool. See `Ice.ThreadPool.name.Size` for more information.

adapter.ThreadPool.SizeMax

Synopsis

`adapter.ThreadPool.SizeMax=num`

Description

num is the maximum number of threads for the thread pool. See `Ice.ThreadPool.name.SizeMax` for more information.

The default value is the value of `adapter.ThreadPool.Size`, meaning the thread pool can never grow larger than its initial size.

adapter.ThreadPool.SizeWarn

Synopsis

`adapter.ThreadPool.SizeWarn=num`

Description

Whenever *num* threads are active in a thread pool, a "low on threads" warning is printed. The default value is zero, which disables the warning.

adapter.ThreadPool.StackSize

Synopsis

`adapter.ThreadPool.StackSize=num`

Description

num is the stack size (in bytes) of threads in the thread pool. The default value is zero, meaning the operating system's default is used.

adapter.ThreadPool.ThreadIdleTime

Synopsis

Copyright 2016, ZeroC, Inc.

adapter.ThreadPool.ThreadIdleTime=*num*

Description

In a dynamically-sized thread pool, Ice reaps a thread after it is idle for *num* seconds. Setting this property to zero disables idle thread reaping. If not specified, the default value is 60 seconds. See `Ice.ThreadPool.`*name*`.ThreadIdleTime` for more information.

adapter.ThreadPool.ThreadPriority

Synopsis

adapter.ThreadPool.ThreadPriority=*num*

Description

num specifes a thread priority for the object adapter's thread pool. The object adapter creates its threads with the specified priority. Leaving this property unset causes the adapter to create threads with the priority specified by `Ice.ThreadPool.Server.ThreadPriority` or, if that property is unset, the priority specified by `Ice.ThreadPriority`.

Copyright 2016, ZeroC, Inc.

Proxy Properties

The communicator operation `propertyToProxy` creates a proxy from a group of configuration properties. The argument to `propertyToProxy` is a string representing the base name of the property group (shown as *name* in the property descriptions below). This name must correspond to a property that supplies the stringified form of the proxy. Subordinate properties can be defined to customize the proxy's local configuration.

The communicator operation `proxyToProperty` performs the inverse operation, that is, returns the property group for a proxy.

On this page:

- name
- name.CollocationOptimized
- name.ConnectionCached
- name.Context.key
- name.EndpointSelection
- name.InvocationTimeout
- name.Locator
- name.LocatorCacheTimeout
- name.PreferSecure
- name.Router

name

Synopsis

name=*proxy*

Description

The base property of the group with an application-specific *name* supplying the stringified representation of a proxy. Use the communicator operation `propertyToProxy` to retrieve the property and convert it into a proxy.

name.CollocationOptimized

Synopsis

name.CollocationOptimized=*num*

Description

If *num* is a value greater than zero, the proxy is configured to use collocated invocations when possible. Defining this property is equivalent to invoking the `ice_collocationOptimized` proxy method.

name.ConnectionCached

Synopsis

name.ConnectionCached=*num*

Description

If *num* is a value greater than zero, the proxy caches its chosen connection for use in subsequent requests. Defining this property is equivalent to invoking the `ice_connectionCached` proxy method.

name.Context.*key*

Copyright 2016, ZeroC, Inc.

Synopsis

name.Context.*key*=*value*

Description

Adds the key/value pair to the proxy's request context.

name.EndpointSelection

Synopsis

name.EndpointSelection=*type*

Description

Specifies the proxy's endpoint selection type. Legal values are `Random` and `Ordered`. Defining this property is equivalent to invoking the `ice_endpointSelection` proxy method.

name.InvocationTimeout

Synopsis

name.InvocationTimeout=*num*

Description

Specifies the invocation timeout to be used by this proxy. Defining this property is equivalent to invoking the `ice_invocationTimeout` proxy method.

name.Locator

Synopsis

name.Locator=*proxy*

Description

Specifies the proxy of the locator to be used by this proxy. Defining this property is equivalent to invoking the `ice_locator` proxy method.

This is a proxy property, so you can configure additional local aspects of the proxy with subordinate properties. For example:

```
MyProxy.Locator=...
MyProxy.Locator.EndpointSelection=Ordered
```

name.LocatorCacheTimeout

Synopsis

name.LocatorCacheTimeout=*num*

Description

Copyright 2016, ZeroC, Inc.

Specifies the locator cache timeout to be used by this proxy. Defining this property is equivalent to invoking the `ice_locatorCacheTimeout` proxy method.

name.PreferSecure

Synopsis

name`.PreferSecure=`*num*

Description

If *num* is a value greater than zero, the proxy gives precedence to secure endpoints. If not defined, the proxy uses the value of `Ice.Default.PreferSecure`.

Defining this property is equivalent to invoking the `ice_preferSecure` proxy method.

name.Router

Synopsis

name`.Router=`*proxy*

Description

Specifies the proxy of the router to be used by this proxy. Defining this property is equivalent to invoking the `ice_router` proxy method.

This is a proxy property, so you can configure additional local aspects of the proxy with subordinate properties. For example:

```
MyProxy.Router=...
MyProxy.Router.PreferSecure=1
```


Copyright 2016, ZeroC, Inc.

Miscellaneous Ice.* Properties

On this page:

- Ice.BackgroundLocatorCacheUpdates
- Ice.BatchAutoFlushSize
- Ice.CacheMessageBuffers
- Ice.ChangeUser
- Ice.CollectObjects
- Ice.Compression.Level
- Ice.ConsoleListener
- Ice.EventLog.Source
- Ice.HTTPProxyHost
- Ice.HTTPProxyPort
- Ice.ImplicitContext
- Ice.LogFile
- Ice.LogFile.SizeMax
- Ice.LogStdErr.Convert
- Ice.MessageSizeMax
- Ice.Nohup
- Ice.NullHandleAbort
- Ice.Package.module
- Ice.PrintAdapterReady
- Ice.PrintProcessId
- Ice.PrintStackTraces
- Ice.ProgramName
- Ice.RetryIntervals
- Ice.ServerIdleTime
- Ice.SOCKSProxyHost
- Ice.SOCKSProxyPort
- Ice.StdErr
- Ice.StdOut
- Ice.SyslogFacility
- Ice.ThreadInterruptSafe
- Ice.UseSyslog
- Ice.Voip

Ice.BackgroundLocatorCacheUpdates

Synopsis

`Ice.BackgroundLocatorCacheUpdates=`*`num`*

Description

If *`num`* is set to zero (the default), an invocation on an indirect proxy whose endpoints are older than the configured locator cache timeout triggers a locator cache update; the run time delays the invocation until the new endpoints are returned by the locator.

If *`num`* is set to a value larger than zero, an invocation on an indirect proxy with expired endpoints still triggers a locator cache update, but the update is performed in the background, and the run time uses the expired endpoints for the invocation. This avoids delaying the first invocation that follows expiry of a cache entry.

Ice.BatchAutoFlushSize

Synopsis

`Ice.BatchAutoFlushSize=`*`num`*

Description

This property controls how the Ice run time deals with flushing of batch messages. If *`num`* is set to a value greater than zero, the run time automatically forces a flush of the current batch when a new message is added to a batch and that message would cause the batch to exceed *`num`* kilobytes. If *`num`* is set to zero or a negative number, batches must be flushed explicitly by the application. If not defined, the default value is `1024`.

Copyright 2016, ZeroC, Inc.

When flushed, batch requests are sent as a single Ice message. The Ice run time in the receiver limits incoming messages to the maximum size specified by `Ice.MessageSizeMax`, therefore the sender must periodically flush batch requests (whether manually or automatically) to ensure they do not exceed the receiver's configured limit.

Ice.CacheMessageBuffers

Synopsis

`Ice.CacheMessageBuffers=`*`num`* (Java, .NET)

Description

If *`num`* is a value greater than zero, the Ice run time caches message buffers for future reuse. This can improve performance and reduce the amount of garbage produced by Ice internals that the garbage collector would eventually spend time to reclaim. However, for applications that exchange very large messages, this cache may consume excessive amounts of memory and therefore should be disabled by setting this property to zero.

This property affects the caching of message buffers for synchronous invocation, asynchronous invocation, and synchronous dispatch. The Ice run time never caches message buffers for asynchronous dispatch.

Platform Notes

- Java
  Ice allocates non-direct message buffers when this property is set to 1 and direct message buffers when set to 2. Use of direct message buffers minimizes copying and typically results in improved throughput. If not defined, the default value is 2.
- .NET
  If not defined, the default value is 1.

Ice.ChangeUser

Synopsis

`Ice.ChangeUser=`*`user`* (C++ & Unix only)

Description

If set, Ice changes the user and group id to the respective ids of *`user`* in `/etc/passwd`. This only works if the Ice application is executed by the super-user.

Ice.CollectObjects

Synopsis

`Ice.CollectObjects=`*`num`* (C++)

Description

Ice for C++ includes a garbage collection facility for reclaiming cyclic graphs of Slice class instances that are unmarshaled by the Ice run time. Setting this property to `1` causes the Ice run time to assume that all cyclic object graphs that a program receives are eligible for collection by default. If not specified, the default value is zero.

Ice.Compression.Level

Synopsis

`Ice.Compression.Level=`*`num`*

Copyright 2016, ZeroC, Inc.

Description

Specifies the bzip2 compression level to use when compressing protocol messages. Legal values for *num* are 1 to 9, where 1 represents the fastest compression and 9 represents the best compression. Note that higher levels cause the bzip2 algorithm to devote more resources to the compression effort, and may not result in a significant improvement over lower levels. If not specified, the default value is 1.

Ice.ConsoleListener

Synopsis

`Ice.ConsoleListener=`*num* (.NET)

Description

If *num* is non-zero, the Ice run time installs a `ConsoleTraceListener` that writes its messages to `stderr`. If *num* is zero, logging is disabled. Note that the setting of `#Ice.LogFile` overrides this property: if `Ice.LogFile` is set, messages are written to the log file regardless of the setting of `Ice.ConsoleListener`.

Ice.EventLog.Source

Synopsis

`Ice.EventLog.Source=`*name* (C++ & Windows only)

Description

Specifies the name of an event log source to be used by a Windows service that subclasses `Ice::Service`. The value of *name* represents a subkey of the `Eventlog` registry key. An application (or administrator) typically prepares the registry key when the service is installed. If no matching registry key is found, Windows logs events in the `Application` log. Any backslashes in *name* are silently converted to forward slashes. If not defined, `Ice::Service` uses the service name as specified by the `--service` option.

Ice.HTTPProxyHost

Synopsis

`Ice.HTTPProxyHost=`*addr*

Description

Specifies the host name or IP address of an HTTP proxy server. If *addr* is not empty, Ice uses the designated HTTP proxy server for all outgoing (client) connections.

Ice.HTTPProxyPort

Synopsis

`Ice.HTTPProxyPort=`*num*

Description

The port number of the HTTP proxy server. If not specified, the default value is `1080`.

Ice.ImplicitContext

Synopsis

`Ice.ImplicitContext=`*type*

Copyright 2016, ZeroC, Inc.

Description

Specifies whether a communicator has an implicit request context and, if so, at what scope the context applies. Legal values for this property are `None` (equivalent to the empty string), `PerThread`, and `Shared`. If not specified, the default value is `None`.

Ice.LogFile

Synopsis

`Ice.LogFile=`*`file`*

Description

Replaces the communicator's default logger with a simple file-based logger implementation. This property does not affect the per-process logger. The logger creates the specified file if necessary, otherwise it appends to the file. If the logger is unable to open the file, the application receives an `InitializationException` during communicator initialization. If a logger object is supplied in the `InitializationData` argument during communicator initialization, it takes precedence over this property.

Ice.LogFile.SizeMax

Synopsis

`Ice.LogFile.SizeMax=`*`num`* `(C++)`

Description

num is a positive integer that represents the maximum size of log files configured through `Ice.LogFile`, in bytes. When a log file's size reaches *num*, the Ice file-based logger renames this log file to *baselogfilename-datetimestamp.ext* and creates a new log file. The default value for *num* is 0, which means that the log file's size is unlimited. In this case, the Ice file-based logger opens and writes to a single log file.

Ice.LogStdErr.Convert

Synopsis

`Ice.LogStdErr.Convert=`*`num (C++)`*

Description

If *num* is set to a value larger than zero, on Windows, the communicator's default logger converts log messages from the application's narrow string encoding (as defined by the installed narrow string converter, if any) to the Windows console's code page. The default value for this property is 1 when `Ice.StdErr` is not set, and 0 otherwise. This property is read by the first communicator created in a process; it is ignored by other communicators.

Ice.MessageSizeMax

Synopsis

`Ice.MessageSizeMax=`*`num`*

Description

This property controls the maximum size (in kilobytes) of an uncompressed protocol message that is accepted by the Ice run time. The size includes the size of the Ice protocol header. The default size is `1024` (1 megabyte).

The only purpose of this property is to prevent a malicious or defective sender from triggering a large memory allocation in a receiver. If this is not a concern, you can set `Ice.MessageSizeMax` to zero; as of Ice 3.6, setting this property to zero (or to a negative number) disables the message size limit altogether.

If the Ice run time in a receiver encounters an incoming message whose size exceeds the receiver's setting for `Ice.MessageSizeMax`, the run time raises a `MemoryLimitException` and closes the connection. For example, when a client receives an oversized reply message,

Copyright 2016, ZeroC, Inc.

the result of its invocation is a `MemoryLimitException`. When a server receives an oversized request message, the client receives a `ConnectionLostException` (because the server closed the connection) and the server logs a message if `Ice.Warn.Connections` is set.

Ice.Nohup

Synopsis

`Ice.Nohup=`*`num`*

Description

If *`num`* is set to a value larger than zero, the `Application` convenience class (as well as the `Ice::Service` class in C++) ignore `SIGHUP` on Unix and `CTRL_LOGOFF_EVENT` on Windows. As a result, an application that sets `Ice.Nohup` continues to run if the user that started the application logs off. The default value for `Application` is `0`, and the default value for `Ice::Service` is `1`.

Ice.NullHandleAbort

Synopsis

`Ice.NullHandleAbort=`*`num`*

Description

If *`num`* is set to a value larger than zero, invoking an operation using a null smart pointer causes the program to abort immediately instead of raising `IceUtil::NullHandleException`.

Ice.Package.*module*

Synopsis

`Ice.Package.`*`module`*`=`*`package`* (Java)

Description

Ice for Java allows you to customize the packaging of generated code. If you use this feature, the Ice run time requires additional configuration in order to successfully unmarshal exceptions and concrete class types. This property associates a top-level Slice *`module`* with a Java *`package`*. If all top-level modules are generated into the same user-defined package, it is easier to use `Ice.Default.Package` instead.

Ice.PrintAdapterReady

Synopsis

`Ice.PrintAdapterReady=`*`num`*

Description

If *`num`* is set to a value larger than zero, an object adapter prints "*adapter_name* ready" on standard output after activation is complete. This is useful for scripts that need to wait until an object adapter is ready to be used.

Ice.PrintProcessId

Synopsis

`Ice.PrintProcessId=`*`num`*

Description

Copyright 2016, ZeroC, Inc.

If *num* is set to a value larger than zero, the process ID is printed on standard output upon startup.

Ice.PrintStackTraces

Synopsis

`Ice.PrintStackTraces=`*num* (JavaScript, C++)

Description

If *num* is set to a value larger than zero, inserting an exception that derives from `IceUtil::Exception` into a logger helper class (such as `Ice::Warning`) also displays the exception's stack trace. Likewise, the `ice_stackTrace` function on the base exception class, `IceUtil::Exception`, will return the stack trace or an empty string depending on the value *num*. If not set, the default value depends on how the Ice run time is compiled: 0 for an optimized build and 1 for a debug build.

On Windows, stack traces are available in debug builds and in release builds built with the environment variable `RELEASEPDBS=yes`. The release DLLs included in the standard Ice binary distribution are built with this setting enabled. The `.PDB` files for both the debug and release distribution builds are available in a separate Ice PDBs Windows installer.

When generating a stack trace, Windows will attempt to locate the `.PDB` files at the location stored in the associated DLLs; in case the `.PDB` file is not found at this location, Windows will attempt to locate this file using the following search path: current working directory, then the directory of the `IceUtil` DLL, and finally the search path specified by the environment variable `_NT_SYMBOL_PATH`. Therefore, if your `.PDB` files are no longer where you built them, you will need to copy them next to the `IceUtil` DLL or include their new directory in the `_NT_SYMBOL_PATH` environment variable.

This property is only supported in JavaScript, C++, and the scripting languages that use the Ice for C++ run time (Python, Ruby, PHP). Be aware that enabling this property in Python, Ruby or PHP will only display C/C++ stack traces.

Ice.ProgramName

Synopsis

`Ice.ProgramName=`*name*

Description

name is the program name, which is set automatically from `argv[0]` (C++) and from `AppDomain.CurrentDomain.FriendlyName` (.NET) during initialization. For Java, `Ice.ProgramName` is initialized to the empty string. The default name can be overridden by setting this property.

Ice.RetryIntervals

Synopsis

`Ice.RetryIntervals=`*num* `[`*num* `...]`

Description

This property defines the number of times an operation is automatically retried and the delay between each retry. For example, if the property is set to `0 100 500`, the operation is retried 3 times: immediately after the first failure, again after waiting 100ms after the second failure, and again after waiting 500ms after the third failure. The default value (`0`) means Ice retries once immediately. If set to `-1`, no retry occurs.

Ice.ServerIdleTime

Synopsis

`Ice.ServerIdleTime=`*num*

Copyright 2016, ZeroC, Inc.

Description

If *num* is set to a value larger than zero, Ice automatically calls `Communicator::shutdown` once the communicator has been idle for *num* seconds. This shuts down the communicator's server side and causes all threads waiting in `Communicator::waitForShutdown` to return. After that, a server will typically do some clean-up work before exiting. The default value is zero, meaning that the server will not shut down automatically. This property is often used for servers that are automatically activated by IceGrid.

Ice.SOCKSProxyHost

Synopsis

`Ice.SOCKSProxyHost=`*addr*

Description

Specifies the host name or IP address of a SOCKS proxy server. If *addr* is not empty, Ice uses the designated SOCKS proxy server for all outgoing (client) connections.

Ice currently only supports the SOCKS4 protocol, which means only IPv4 connections are allowed.

Ice.SOCKSProxyPort

Synopsis

`Ice.SOCKSProxyPort=`*num*

Description

The port number of the SOCKS proxy server. If not specified, the default value is `1080`.

Ice.StdErr

Synopsis

`Ice.StdErr=`*filename*

Description

If *filename* is not empty, the standard error stream of this process is redirected to this file, in append mode. This property is checked only for the first communicator that is created in a process.

Ice.StdOut

Synopsis

`Ice.StdOut=`*filename*

Description

If *filename* is not empty, the standard output stream of this process is redirected to this file, in append mode. This property is checked only for the first communicator created in a process.

Ice.SyslogFacility

Synopsis

Copyright 2016, ZeroC, Inc.

`Ice.SyslogFacility=`*`string`* (Unix only)

Description

This property sets the syslog facility to *`string`*. This property has no effect if `Ice.UseSyslog` is not set.

`string` can be any of syslog facilities: `LOG_AUTH, LOG_AUTHPRIV, LOG_CRON, LOG_DAEMON, LOG_FTP, LOG_KERN, LOG_LOCAL0, LOG_LOCAL1, LOG_LOCAL2, LOG_LOCAL3, LOG_LOCAL4, LOG_LOCAL5, LOG_LOCAL6, LOG_LOCAL7, LOG_LPR, LOG_MAIL, LOG_NEWS, LOG_SYSLOG, LOG_USER, LOG_UUCP`.

The default value is `LOG_USER`.

Ice.ThreadInterruptSafe

Synopsis

`Ice.ThreadInterruptSafe=`*`num`* (Java)

Description

If *`num`* is set to a value larger than zero, Ice for Java disables message caching by setting `Ice.CacheMessageBuffers` to 0 and takes other steps necessary to ensure that Java interrupts work correctly. If not defined, the default value is zero.

Ice.UseSyslog

Synopsis

`Ice.UseSyslog=`*`num`* (Unix only)

Description

If *`num`* is set to a value larger than zero, a special logger is installed that logs to the `syslog` service instead of standard error. The identifier for `syslog` is the value of `#Ice.ProgramName`. Use `#Ice.SyslogFacility` to select a `syslog` facility.

Ice.Voip

Synopsis

`Ice.Voip=`*`num`* (Ice Touch only)

Description

If *`num`* is set to a value larger than zero, the Ice run time sets the `kCFStreamNetworkServiceType` property to `kCFStreamNetworkServiceTypeVoIP` for all sockets created through the Ice communicator. Only voice over IP applications are allowed to use this setting. It ensures the sockets aren't closed and enables receiving data when the application is in the background. See the description from Configuring Sockets for VoIP Usage for information on this setting and when it's applicable.

The default value is zero.

Copyright 2016, ZeroC, Inc.

Freeze.*

On this page:

- Freeze.DbEnv.env-name.CheckpointPeriod
- Freeze.DbEnv.env-name.DbHome
- Freeze.DbEnv.env-name.DbPrivate
- Freeze.DbEnv.env-name.DbRecoverFatal
- Freeze.DbEnv.env-name.EncodingVersion
- Freeze.DbEnv.env-name.LockFile
- Freeze.DbEnv.env-name.OldLogsAutoDelete
- Freeze.DbEnv.env-name.PeriodicCheckpointMinSize
- Freeze.Evictor.env-name.filename.MaxTxSize
- Freeze.Evictor.env-name.filename.name.BtreeMinKey
- Freeze.Evictor.env-name.filename.name.Checksum
- Freeze.Evictor.env-name.filename.PageSize
- Freeze.Evictor.env-name.filename.PopulateEmptyIndices
- Freeze.Evictor.env-name.filename.RollbackOnUserException
- Freeze.Evictor.env-name.filename.SavePeriod
- Freeze.Evictor.env-name.filename.SaveSizeTrigger
- Freeze.Evictor.env-name.filename.StreamTimeout
- Freeze.Map.name.BtreeMinKey
- Freeze.Map.name.Checksum
- Freeze.Map.name.PageSize
- Freeze.Trace.DbEnv
- Freeze.Trace.Evictor
- Freeze.Trace.Map
- Freeze.Trace.Transaction
- Freeze.Warn.Deadlocks
- Freeze.Warn.Rollback

Freeze.DbEnv.*env-name*.CheckpointPeriod

Synopsis

`Freeze.DbEnv.`*`env-name`*`.CheckpointPeriod=`*`num`*

Description

Every Berkeley DB environment created by Freeze has an associated thread that checkpoints this environment every *`num`* seconds. If *`num`* is less than `0`, no checkpointing is performed. The default is 120 seconds.

Freeze.DbEnv.*env-name*.DbHome

Synopsis

`Freeze.DbEnv.`*`env-name`*`.DbHome=`*`db-home`*

Description

Defines the home directory of this Freeze database environment. The default directory is *`env-name`*.

Freeze.DbEnv.*env-name*.DbPrivate

Synopsis

`Freeze.DbEnv.`*`env-name`*`.DbPrivate=`*`num`*

Description

Copyright 2016, ZeroC, Inc.

If *num* is set to a value larger than zero, Freeze instructs Berkeley DB to use process-private memory instead of shared memory. The default value is `1`. Set this property to `0` in order to run a FreezeScript utility, or a Berkeley DB utility such as `db_archive`, on a running environment.

Freeze.DbEnv.*env-name*.DbRecoverFatal

Synopsis

`Freeze.DbEnv.`*`env-name`*`.DbRecoverFatal=`*`num`*

Description

If *num* is set to a value larger than zero, fatal recovery is performed when the environment is opened. The default value is `0`.

Freeze.DbEnv.*env-name*.EncodingVersion

Synopsis

`Freeze.DbEnv.`*`env-name`*`.EncodingVersion=`*`encoding`*

Description

Defines the encoding used to decode keys and to encode keys and values. The default value is the value of `Ice.Default.EncodingVersion`.

Freeze.DbEnv.*env-name*.LockFile

Synopsis

`Freeze.DbEnv.`*`env-name`*`.LockFile=`*`num`*

Description

If *num* is set to a value larger than zero, Freeze creates a lock file in the database environment to prevent other processes from opening the environment. The default value is `1`.

Note that applications should not normally disable the lock file because simultaneous access to the same environment by multiple processes can lead to data corruption.

FreezeScript utilities automatically disable the lock file when `Freeze.DbEnv.`*`env-name`*`.DbPrivate` is set to zero.

Freeze.DbEnv.*env-name*.OldLogsAutoDelete

`Freeze.DbEnv.`*`env-name`*`.OldLogsAutoDelete=`*`num`*

If *num* is set to a value larger than zero, old transactional logs no longer in use are deleted after each periodic checkpoint (see `Freeze.DbEnv.`*`env-name`*`.CheckpointPeriod`). The default value is `1`.

Freeze.DbEnv.*env-name*.PeriodicCheckpointMinSize

Synopsis

`Freeze.DbEnv.`*`env-name`*`.PeriodicCheckpointMinSize=`*`num`*

Description

Copyright 2016, ZeroC, Inc.

num is the minimum size in kilobytes for the periodic checkpoint (see `Freeze.DbEnv.env-name.CheckpointPeriod`). This value is passed to Berkeley DB's `checkpoint` function. The default is `0` (which means no minimum).

Freeze.Evictor.*env-name.filename*.MaxTxSize

Synopsis

`Freeze.Evictor.env-name.filename.MaxTxSize=num`

Description

Freeze can use a background thread to save updates to the database. Transactions are used to save many facets together. *num* defines the maximum number of facets saved per transaction. The default is `10 * SaveSizeTrigger` (see `Freeze.Evictor.env-name.filename.SaveSizeTrigger`); if this value is negative, the actual value is set to `100`.

Freeze.Evictor.*env-name.filename.name*.BtreeMinKey

Synopsis

`Freeze.Evictor.env-name.filename.name.BtreeMinKey=num`

Description

name represents a database name or an index name. This property sets the B-tree minkey of the corresponding Berkeley DB database. *num* is ignored if it is less than 2. Please refer to the Berkeley DB documentation for details.

Freeze.Evictor.*env-name.filename.name*.Checksum

Synopsis

{{Freeze.Evictor.*env-name.filename*.Checksum=*num*

Description

If *num* is greater than 0, checksums on the corresponding Berkeley DB database(s) are enabled. Please refer to the Berkeley DB documentation for details.

Freeze.Evictor.*env-name.filename*.PageSize

Synopsis

`Freeze.Evictor.env-name.filename.PageSize=num`

Description

If *num* is greater than 0, it sets the page size of the corresponding Berkeley DB database(s). Please refer to the Berkeley DB documentation f or details.

Freeze.Evictor.*env-name.filename*.PopulateEmptyIndices

Synopsis

`Freeze.Evictor.env-name.filename.PopulateEmptyIndices=num`

Description

Copyright 2016, ZeroC, Inc.

When *num* is not `0` and you create an evictor with one or more empty indexes, the `createBackgroundSaveEvictor` or `createTransactionalEvictor` call will populate these indexes by iterating over all the corresponding facets. This is particularly useful after transforming a Freeze evictor with FreezeScript, since FreezeScript does not transform indexes; however this can significantly slow down the creation of the evictor if you have an empty index because none of the facets currently in the database match the type of this index. The default value for this property is `0`.

Freeze.Evictor.*env-name.filename*.RollbackOnUserException

Synopsis

`Freeze.Evictor.`*`env-name`*`.`*`filename`*`.RollbackOnUserException=`*`num`*

Description

If *num* is set to a value larger than zero, a transactional evictor rolls back its transaction if the outcome of the dispatch is a user exception. If *num* is 0 (the default), the transactional evictor commits the transaction.

Freeze.Evictor.*env-name.filename*.SavePeriod

Synopsis

`Freeze.Evictor.`*`env-name`*`.`*`filename`*`.SavePeriod=`*`num`*

Description

Freeze can use a background thread to save updates to the database. After *num* milliseconds without saving, if any facet is created, modified, or destroyed, this background thread wakes up to save these facets. When num is `0`, there is no periodic saving. The default is `60000`.

Freeze.Evictor.*env-name.filename*.SaveSizeTrigger

Synopsis

`Freeze.Evictor.`*`env-name`*`.`*`filename`*`.SaveSizeTrigger=`*`num`*

Description

Freeze can use a background thread to save updates to the database. When *num* is `0` or positive, as soon as *num* or more facets have been created, modified, or destroyed, this background thread wakes up to save them. When *num* is negative, there is no size trigger. The default is `10`.

Freeze.Evictor.*env-name.filename*.StreamTimeout

Synopsis

`Freeze.Evictor.`*`env-name`*`.`*`filename`*`.StreamTimeout=`*`num`*

Description

When the saving thread saves an object, it needs to lock this object in order to get a consistent copy of the object's state. If the lock cannot be acquired within *num* seconds, a fatal error is generated. If a fatal error callback was registered by the application, this callback is called; otherwise the program is terminated immediately. When *num* is `0` or negative, there is no timeout. The default value is `0`.

Freeze.Map.*name*.BtreeMinKey

Copyright 2016, ZeroC, Inc.

Synopsis

`Freeze.Map.`*`name`*`.BtreeMinKey=`*`num`*

Description

`name` may represent a database name or an index name. This property sets the B-tree minkey of the corresponding Berkeley DB database. *`num`* is ignored if it is less than 2. Please refer to the Berkeley DB documentation for details.

Freeze.Map.*name*.Checksum

Synopsis

`Freeze.Map.`*`name`*`.Checksum=`*`num`*

Description

`name` may represent a database name or an index name. If *`num`* is greater than 0, checksums for the corresponding Berkeley DB database are enabled. Please refer to the Berkeley DB documentation for details.

Freeze.Map.*name*.PageSize

Synopsis

`Freeze.Map.`*`name`*`.PageSize=`*`num`*

Description

`name` may represent a database name or an index name. If *`num`* is greater than 0, it sets the page size of the corresponding Berkeley DB database. Please refer to the Berkeley DB documentation for details.

Freeze.Trace.DbEnv

Synopsis

`Freeze.Trace.DbEnv=`*`num`*

Description

The Freeze database environment activity trace level:

0	No database environment activity trace (default).
1	Trace database open and close.
2	Also trace checkpoints and the removal of old log files.

Freeze.Trace.Evictor

Synopsis

`Freeze.Trace.Evictor=`*`num`*

Description

The Freeze evictor activity trace level:

Copyright 2016, ZeroC, Inc.

0	No evictor activity trace (default).
1	Trace Ice object and facet creation and destruction, facet streaming time, facet saving time, object eviction (every 50 objects) and evictor deactivation.
2	Also trace object lookups, and all object evictions.
3	Also trace object retrieval from the database.

Freeze.Trace.Map

Synopsis

`Freeze.Trace.Map=`*`num`*

Description

The Freeze map activity trace level:

0	No map activity trace (default).
1	Trace database open and close.
2	Also trace iterator and transaction operations, and reference counting of the underlying database.

Freeze.Trace.Transaction

Synopsis

`Freeze.Trace.Transaction=`*`num`*

Description

The Freeze transaction activity trace level:

0	No transaction activity trace (default).
1	Trace transaction IDs and commit and rollback activity.

Freeze.Warn.Deadlocks

Synopsis

`Freeze.Warn.Deadlocks=`*`num`*

Description

If *`num`* is set to a value larger than zero, Freeze logs a warning message when a deadlock occur. The default value is `0`.

Freeze.Warn.Rollback

Synopsis

`Freeze.Warn.Deadlocks=`*`num`*

Description

If *`num`* is set to a value larger than zero, Freeze logs a warning message when it rolls back a transaction that goes out of scope together with its associated connection. The default value is `1`. (C++ only)

Copyright 2016, ZeroC, Inc.

Copyright 2016, ZeroC, Inc.

Glacier2.*

On this page:

- Glacier2.AddConnectionContext
- Glacier2.AddUserToAllowCategories
- Glacier2.Admin.AdapterProperty
- Glacier2.AllowCategories
- Glacier2.Client.AdapterProperty
- Glacier2.Client.AlwaysBatch
- Glacier2.Client.Buffered
- Glacier2.Client.ForwardContext
- Glacier2.Client.SleepTime
- Glacier2.Client.Trace.Override
- Glacier2.Client.Trace.Reject
- Glacier2.Client.Trace.Request
- Glacier2.CryptPasswords
- Glacier2.Filter.AdapterId.Accept
- Glacier2.Filter.Address.Accept
- Glacier2.Filter.Address.Reject
- Glacier2.Filter.Category.Accept
- Glacier2.Filter.Category.AcceptUser
- Glacier2.Filter.Identity.Accept
- Glacier2.Filter.ProxySizeMax
- Glacier2.InstanceName
- Glacier2.PermissionsVerifier
- Glacier2.ReturnClientProxy
- Glacier2.RoutingTable.MaxSize
- Glacier2.Server.AdapterProperty
- Glacier2.Server.AlwaysBatch
- Glacier2.Server.Buffered
- Glacier2.Server.ForwardContext
- Glacier2.Server.SleepTime
- Glacier2.Server.Trace.Override
- Glacier2.Server.Trace.Request
- Glacier2.SessionManager
- Glacier2.SessionTimeout
- Glacier2.SSLPermissionsVerifier
- Glacier2.SSLSessionManager
- Glacier2.Trace.RoutingTable
- Glacier2.Trace.Session

Glacier2.AddConnectionContext

Synopsis

`Glacier2.AddConnectionContext=`*`num`*

Description

If *`num`* is set to 1 or 2, Glacier2 adds a number of key-value pairs to the request context that it sends with each request. If *`num`* is set to 1, these entries are added to the context for all forwarded requests. If *`num`* is set to 2, the contexts are added only to calls to `checkPermissions` and `authorize` on permission verifiers, and to calls to `create` on session managers.

If num is non-0, Glacier2 adds the following context entries:

`_con.type`	The type of the connection as returned by `Connection::type`.
`_con.localAddress`	The local address (TCP and SSL only).
`_con.localPort`	The local port (TCP and SSL only).
`_con.remoteAddress`	The remote address (TCP and SSL only).
`_con.remotePort`	The remote port (TCP and SSL only).

Copyright 2016, ZeroC, Inc.

`_con.cipher`	The cipher (SSL only).
`_con.peerCert`	The first certificate of the client certificate chain (SSL only).

The default value is 0, meaning that no contexts are added.

Glacier2.AddUserToAllowCategories

Synopsis

`Glacier2.AddUserToAllowCategories=`*`num`*

Description

Specifies whether to add an authenticated user ID to the `Glacier2.AllowCategories` property when creating a new session. The legal values are shown below:

0	Do not add the user ID (default).
1	Add the user ID.
2	Add the user ID with a leading underscore.

This property is deprecated and supported only for backward-compatibility. New applications should use `Glacier2.Filter.Category.AcceptUser`.

Glacier2.Admin.*AdapterProperty*

Synopsis

`Glacier2.Admin.`*`AdapterProperty=value`*

Description

Glacier2 uses the adapter name `Glacier2.Admin` for its administrative object adapter. Therefore, adapter properties can be used to configure this adapter.

The `Glacier2.Admin.Endpoints` property must be defined to enable the administrative object adapter.

Glacier2's administrative interface allows a remote client to shut down the router; we generally recommend the use of endpoints that are accessible only from behind a firewall.

Glacier2.AllowCategories

Synopsis

`Glacier2.AllowCategories=`*`list`*

Description

Specifies a white space-separated list of identity categories. If this property is defined, then the Glacier2 router only allows requests to Ice objects with an identity that matches one of the categories from this list. If `Glacier2.AddUserToAllowCategories` is defined with a non-0 value, the router automatically adds the user ID of each session to this list.

This property is deprecated and supported only for backward-compatibility. New applications should use `Glacier2.Filter.Category.Accept`.

Glacier2.Client.*AdapterProperty*

Copyright 2016, ZeroC, Inc.

Synopsis

```
Glacier2.Client.AdapterProperty=value
```

Description

Glacier2 uses the adapter name `Glacier2.Client` for the object adapter that it provides to clients. Therefore, adapter properties can be used to configure this adapter.

This adapter must be accessible to clients of Glacier2. Use of a secure transport for this adapter is highly recommended.

Note that `Glacier2.Registry.Client.Endpoints` controls the client endpoint for Glacier2. The port numbers 4063 (for TCP) and 4064 (for SSL) are reserved for Glacier2 by the Internet Assigned Numbers Authority (IANA).

Glacier2.Client.AlwaysBatch

Synopsis

```
Glacier2.Client.AlwaysBatch=num
```

Description

If *num* is set to a value larger than 0, the Glacier2 router always batches queued oneway requests from clients to servers regardless of the value of their `_fwd` contexts. This property is only relevant when `Glacier2.Client.Buffered` is enabled. The default value is `0`.

Glacier2.Client.Buffered

Synopsis

```
Glacier2.Client.Buffered=num
```

Description

If *num* is set to a value larger than 0, the Glacier2 router operates in buffered mode, in which incoming requests from clients are queued and processed in a separate thread. If *num* is set to 0, the router operates in unbuffered mode in which a request is forwarded in the same thread that received it. The default value is `1`.

Glacier2.Client.ForwardContext

Synopsis

```
Glacier2.Client.ForwardContext=num
```

Description

If *num* is set to a value larger than 0, the Glacier2 router includes the request context when forwarding requests from clients to servers. The default value is `0`.

Glacier2.Client.SleepTime

Synopsis

```
Glacier2.Client.SleepTime=num
```

Description

If *num* is set to a value larger than 0, the Glacier2 router sleeps for the specified number of milliseconds after forwarding all queued requests from a client. This delay is useful for batched delivery because it makes it more likely for events to accumulate in a single batch. Similarly, if overrides are used, the delay makes it more likely for overrides to actually take effect. This property is only relevant when `Glacier2.Clien`

Copyright 2016, ZeroC, Inc.

`t.Buffered` is enabled. The default value is `0`.

Glacier2.Client.Trace.Override

Synopsis

`Glacier2.Client.Trace.Override=`*`num`*

Description

If *`num`* is set to a value larger than 0, the Glacier2 router logs a trace message whenever a request was overridden. The default value is `0`.

Glacier2.Client.Trace.Reject

Synopsis

`Glacier2.Client.Trace.Reject=`*`num`*

Description

If *`num`* is set to a value larger than 0, the Glacier2 router logs a trace message whenever the router's configured filters reject a client's request. The default value is `0`.

Glacier2.Client.Trace.Request

Synopsis

`Glacier2.Client.Trace.Request=`*`num`*

Description

If *`num`* is set to a value larger than 0, the Glacier2 router logs a trace message for each request that is forwarded from a client. The default value is `0`.

Glacier2.CryptPasswords

Synopsis

`Glacier2.CryptPasswords=`*`file`*

Description

Specifies the file name of a Glacier2 access control list. Each line of the file must contain a user name and an encrypted password, separated by white space, as described in Writing a Password File. This property is ignored if `Glacier2.PermissionsVerifier` is defined.

Glacier2.Filter.AdapterId.Accept

Synopsis

`Glacier2.Filter.AdapterId.Accept=`*`list`*

Description

Specifies a space-separated list of adapter identifiers. If defined, the Glacier2 router filters requests so that it only allows requests to Ice objects with an adapter identifier that matches one of the entries in this list.

Identifiers that contain spaces must be enclosed in single or double quotes. Single or double quotes that appear within an identifier must be escaped with a leading backslash.

Copyright 2016, ZeroC, Inc.

Glacier2.Filter.Address.Accept

Synopsis

`Glacier2.Filter.Address.Accept=`*`list`*

Description

Specifies a space-separated list of address-port pairs. When defined, the Glacier2 router filters requests so that it only allows requests to Ice objects through proxies that contain network endpoint information that matches an address-port pair listed in this property. If not defined, the default value is `*`, which indicates that any network address is permitted. Requests accepted by this property may be rejected by the `Glacier2.Filter.Address.Reject` property.

Each pair is of the form *`address`*`:`*`port`*. The *`address`* or *`port`* number portion can include wildcards ('`*`') or value ranges or groups. Ranges and groups have the form `[`*`value1, value2, value3, ...`*`]` and/or `[`*`value1-value2`*`]`. Wildcards, ranges, and groups may appear anywhere in the address-port pair string.

Glacier2.Filter.Address.Reject

Synopsis

`Glacier2.Filter.Address.Reject=`*`list`*

Description

Specifies a space-separated list of address-port pairs. When defined, the Glacier2 router rejects requests to Ice objects through proxies that contain network endpoint information that matches an address-port pair listed in this property. If not set, the Glacier2 router allows requests to any network address unless the `Glacier2.Filter.Address.Accept` property is set, in which case requests will be accepted or rejected based on the `Glacier2.Filter.Address.Accept` property. If both the `Glacier2.Filter.Address.Accept` and `Glacier2.Filter.Address.Reject` properties are defined, the `Glacier2.Filter.Address.Reject` property takes precedence.

Each pair is of the form *`address`*`:`*`port`*. The *`address`* or *`port`* number portion can include wildcards ('`*`') or value ranges or groups. Ranges and groups have the form `[`*`value1, value2, value3, ...`*`]` and/or `[`*`value1-value2`*`]`. Wildcards, ranges, and groups may appear anywhere in the address-port pair string.

Glacier2.Filter.Category.Accept

Synopsis

`Glacier2.Filter.Category.Accept=`*`list`*

Description

Specifies a space-separated list of identity categories. If defined, the Glacier2 router filters requests so that it only allows requests to Ice objects with an identity that matches one of the categories in this list. If `Glacier2.Filter.Category.AcceptUser` is defined with a non-0 value, the router automatically adds the user name of each session to this list.

Categories that contain spaces must be enclosed in single or double quotes. Single or double quotes that appear within a category must be escaped with a leading backslash.

Glacier2.Filter.Category.AcceptUser

Synopsis

`Glacier2.Filter.Category.AcceptUser=`*`num`*

Description

Specifies whether to add an authenticated user ID to the `Glacier2.Filter.Category.Accept` property when creating a new session. The legal values are shown below:

Copyright 2016, ZeroC, Inc.

0	Do not add the user ID (default).
1	Add the user ID.
2	Add the user ID with a leading underscore.

Glacier2.Filter.Identity.Accept

Synopsis

`Glacier2.Filter.Identity.Accept=`*`list`*

Description

Specifies a space-separated list of identities. If defined, the Glacier2 router filters requests so that it only allows requests to Ice objects with an identity that matches one of the entries in this list.

Identities that contain spaces must be enclosed in single or double quotes. Single or double quotes that appear within an identity must be escaped with a leading backslash.

Glacier2.Filter.ProxySizeMax

Synopsis

`Glacier2.Filter.ProxySizeMax=`*`num`*

Description

If set, the Glacier2 router rejects requests whose stringified proxies are longer than *`num`*. This helps secure the system against attack. If not set, Glacier2 will accept requests using proxies of any length.

Glacier2.InstanceName

Synopsis

`Glacier2.InstanceName=`*`name`*

Description

Specifies a default identity category for the Glacier2 objects. If defined, the identity of the Glacier2 administrative interface becomes *`name`*`/admin` and the identity of the Glacier2 router interface becomes *`name`*`/router`.

If not defined, the default value is `Glacier2`.

Glacier2.PermissionsVerifier

Synopsis

`Glacier2.PermissionsVerifier=`*`proxy`*

Description

Specifies the proxy of an object that implements the `Glacier2::PermissionsVerifier` interface for controlling access to Glacier2 sessions. The router invokes this proxy to validate the user name and password of each new session. Sessions created from a secure connection are verified by the object specified in `Glacier2.SSLPermissionsVerifier`. For simple configurations, you can specify the name of a password file using `Glacier2.CryptPasswords`.

Glacier2 supplies a "null" permissions verifier object that accepts any username and password combination for situations in which no authentication is necessary. To enable this verifier, set the property value to *`instance`*`/NullPermissionsVerifier`, where *`instance`* is the value of `Glacier2.InstanceName`.

Copyright 2016, ZeroC, Inc.

As a proxy property, you can configure additional aspects of the proxy using properties.

Glacier2.ReturnClientProxy

Synopsis

`Glacier2.ReturnClientProxy=`*`num`*

Description

If *`num`* is a value greater than 0, Glacier2 maintains backward compatibility with clients using Ice versions prior to 3.2.0. In this case you should also define `Glacier2.Client.PublishedEndpoints` to specify the endpoints that clients should use to contact the router. For example, if the Glacier2 router resides behind a network firewall, the `Glacier2.Client.PublishedEndpoints` property should specify the firewall's external endpoints.

If not defined, the default value is 0.

Glacier2.RoutingTable.MaxSize

Synopsis

`Glacier2.RoutingTable.MaxSize=`*`num`*

Description

This property sets the size of the router's routing table to *`num`* entries. If more proxies are added to the table than this value, proxies are evicted from the table on a least-recently used basis.

Clients based on Ice version 3.1 and later automatically retry operation calls on evicted proxies and transparently re-add such proxies to the table. Clients based on Ice versions earlier than 3.1 receive an `ObjectNotExistException` for invocations on evicted proxies. For such older clients, *`num`* must be set to a sufficiently large value to prevent these clients from failing.

The default size of the routing table is 1000.

Glacier2.Server.*AdapterProperty*

Synopsis

`Glacier2.Server.`*`AdapterProperty=value`*

Description

Glacier2 uses the adapter name `Glacier2.Server` for the object adapter that it provides to servers. Therefore, adapter properties can be used to configure this adapter.

This adapter provides access to the `SessionControl` interface and must be accessible to servers that call back to router clients.

Glacier2.Server.AlwaysBatch

Synopsis

`Glacier2.Server.AlwaysBatch=`*`num`*

Description

If *`num`* is set to a value larger than 0, the Glacier2 router always batches queued oneway requests from servers to clients regardless of the value of their `_fwd` contexts. This property is only relevant when `Glacier2.Server.Buffered` is enabled. The default value is `0`.

Glacier2.Server.Buffered

Copyright 2016, ZeroC, Inc.

Synopsis

`Glacier2.Server.Buffered=`*`num`*

Description

If *`num`* is set to a value larger than 0, the Glacier2 router operates in buffered mode, in which incoming requests from servers are queued and processed in a separate thread. If *`num`* is set to 0, the router operates in unbuffered mode in which a request is forwarded in the same thread that received it. The default value is `1`.

Glacier2.Server.ForwardContext

Synopsis

`Glacier2.Server.ForwardContext=`*`num`*

Description

If *`num`* is set to a value larger than 0, the Glacier2 router includes the request context when forwarding requests from servers to clients. The default value is `0`.

Glacier2.Server.SleepTime

Synopsis

`Glacier2.Server.SleepTime=`*`num`*

Description

If *`num`* is set to a value larger than 0, the Glacier2 router sleeps for the specified number of milliseconds after forwarding all queued requests from a server. This delay is useful for batched delivery because it makes it more likely for events to accumulate in a single batch. Similarly, if overrides are used, the delay makes it more likely for overrides to actually take effect. This property is only relevant when `Glacier2.Server.Buffered` is enabled. The default value is `0`.

Glacier2.Server.Trace.Override

Synopsis

`Glacier2.Server.Trace.Override=`*`num`*

Description

If *`num`* is set to a value larger than 0, the Glacier2 router logs a trace message whenever a request is overridden. The default value is `0`.

Glacier2.Server.Trace.Request

Synopsis

`Glacier2.Server.Trace.Request=`*`num`*

Description

If *`num`* is set to a value larger than 0, the Glacier2 router logs a trace message for each request that is forwarded from a server. The default value is `0`.

Glacier2.SessionManager

Synopsis

Copyright 2016, ZeroC, Inc.

`Glacier2.SessionManager=`*`proxy`*

Description

Specifies the proxy of an object that implements the `Glacier2::SessionManager` interface. The router invokes this proxy to create a new session for a client, but only after the router validates the client's user name and password.

As a proxy property, you can configure additional aspects of the proxy using properties.

Glacier2.SessionTimeout

Synopsis

`Glacier2.SessionTimeout=`*`num`*

Description

If *`num`* is set to a value larger than 0, a client's session with the Glacier2 router expires after the specified *`num`* seconds of inactivity. The default value is `0`, meaning sessions do not expire due to inactivity.

It is important to choose *`num`* such that client sessions do not expire prematurely.

The session timeout may also affect Active Connection Management (ACM) for client connections, which are connections to the router's object adapter named `Glacier2.Client`. If you have not explicitly configured the router with a value for `Glacier2.Client.ACM.Timeout`, the router uses the session timeout as the value for this property. If no session timeout is defined, the router's incoming client connections use its default ACM behavior.

Glacier2.SSLPermissionsVerifier

Synopsis

`Glacier2.SSLPermissionsVerifier=`*`proxy`*

Description

Specifies the proxy of an object that implements the `Glacier2::SSLPermissionsVerifier` interface for controlling access to Glacier2 sessions. The router invokes this proxy to verify the credentials of clients that attempt to create a session from a secure connection. Sessions created with a user name and password are verified by the object specified in `Glacier2.PermissionsVerifier`.

Glacier2 supplies a "null" permissions verifier object that accepts the credentials of any client for situations in which no authentication is necessary. To enable this verifier, set the property value to *`instance`*`/NullSSLPermissionsVerifier`, where *`instance`* is the value of `Glacier2.InstanceName`.

As a proxy property, you can configure additional aspects of the proxy using properties.

Glacier2.SSLSessionManager

Synopsis

`Glacier2.SSLSessionManager=`*`proxy`*

Description

Specifies the proxy of an object that implements the `Glacier2::SSLSessionManager` interface for managing sessions. The router invokes this proxy to create a new session for a client that has called `createSessionFromSecureConnection`.

As a proxy property, you can configure additional aspects of the proxy using properties.

Glacier2.Trace.RoutingTable

Synopsis

Copyright 2016, ZeroC, Inc.

`Glacier2.Trace.RoutingTable=`*`num`*

Description

The routing table trace level:

0	No routing table trace (default).
1	Logs a message for each proxy that is added to the routing table.
2	Logs a message for each proxy that is evicted from the routing table (see `#Glacier2.RoutingTable.MaxSize`).
3	Combines the output for trace levels 1 and 2.

Glacier2.Trace.Session

Synopsis

`Glacier2.Trace.Session=`*`num`*

Description

If *`num`* is set to a value larger than 0, the Glacier2 router logs trace messages about session-related activities. The default value is `0`.

Copyright 2016, ZeroC, Inc.

Ice.ACM.*

On this page:

- Ice.ACM.Client.Close
- Ice.ACM.Client.Heartbeat
- Ice.ACM.Client.Timeout
- Ice.ACM.Close
- Ice.ACM.Heartbeat
- Ice.ACM.Timeout
- Ice.ACM.Server.Close
- Ice.ACM.Server.Heartbeat
- Ice.ACM.Server.Timeout

Ice.ACM.Client.Close

Synopsis

`Ice.ACM.Client.Close=`*`num`*

Description

Overrides the value of `Ice.ACM.Close` in a client context, meaning for outgoing connections.

Ice.ACM.Client.Heartbeat

Synopsis

`Ice.ACM.Client.Heartbeat=`*`num`*

Description

Overrides the value of `Ice.ACM.Heartbeat` in a client context, meaning for outgoing connections.

Ice.ACM.Client.Timeout

Synopsis

`Ice.ACM.Client.Timeout=`*`num`*

Description

Overrides the value of `Ice.ACM.Timeout` in a client context, meaning for outgoing connections.

Ice.ACM.Close

Synopsis

`Ice.ACM.Close=`*`num`*

Description

This property determines when Active Connection Management (ACM) closes a connection. Legal values for *`num`* correspond directly to the `ACMClose` enumerators found in `Connection.ice`:

Copyright 2016, ZeroC, Inc.

Slice

```
local enum ACMClose {
    CloseOff,
    CloseOnIdle,
    CloseOnInvocation,
    CloseOnInvocationAndIdle,
    CloseOnIdleForceful
};
```

The table below describes the semantics of each value:

Enumerator	Property Value	Description
`CloseOff`	0	Disables automatic connection closure. A connection will be closed when the communicator is destroyed, when there's a network failure or when the peer closes it.
`CloseOnIdle`	1	Gracefully closes a connection that has been idle for the configured timeout period. A connection is idle if it has no pending outgoing or incoming requests.
`CloseOnInvocation`	2	Forcefully closes a connection that has been idle for the configured timeout period, but only if the connection has pending outgoing requests. This is useful when you don't want a connection to be closed when idle but still want the connection to be forcefully closed when invocations are in progress and no messages are received from the server (potentially indicating that the server is dead). This should be used with a heartbeat configuration where heartbeats are sent at regular intervals by a server while an invocation is being dispatched.
`CloseOnInvocationAndIdle`	3	Combines the behaviors of `CloseOnIdle` and `CloseOnInvocation`.
`CloseOnIdleForceful`	4	Forcefully closes a connection that has been idle for the configured timeout period, regardless of whether the connection has pending outgoing or incoming requests. This is typically used together with a heartbeat configuration that keeps idle connections alive.

You must set a non-zero value for `Ice.ACM.Timeout` for this property to have any effect. Use `Ice.ACM.Close` in conjunction with `Ice.ACM.Heartbeat` to tailor a connection management policy for your application. Use the corresponding `Ice.ACM.Client.Close` or `Ice.ACM.Server.Close` properties to override this setting in a client or server context, respectively. See Connection Closure for more information on closing connections.

If not defined, the default value is 3 in a client context and 2 in a server context.

Ice.ACM.Heartbeat

Synopsis

`Ice.ACM.Heartbeat=`*`num`*

Description

This property determines whether heartbeats are enabled. You can use heartbeats in conjunction with the `Ice.ACM.Close` and `Ice.ACM.Timeout` properties to create a connection management policy as part of Active Connection Management (ACM). Legal values for *`num`* correspond directly to the `ACMHeartbeat` enumerators found in `Connection.ice`:

Copyright 2016, ZeroC, Inc.

Slice

```
local enum ACMHeartbeat {
    HeartbeatOff,
    HeartbeatOnInvocation,
    HeartbeatOnIdle,
    HeartbeatAlways
};
```

The table below describes the semantics of each value:

Enumerator	Property Value	Description
`HeartbeatOff`	0	Disables client-side heartbeats.
`HeartbeatOnInvocation`	1	Send a heartbeat at regular intervals if there are pending incoming requests.
`HeartbeatOnIdle`	2	Send a heartbeat at regular intervals when the connection is idle.
`HeartbeatAlways`	3	Send a heartbeat at regular intervals until the connection is closed.

You must set a non-zero value for `Ice.ACM.Timeout` for this property to have any effect. Use the corresponding `Ice.ACM.Client.Heartbeat` or `Ice.ACM.Server.Heartbeat` properties to override this setting in a client or server context, respectively.

To send a heartbeat, the Ice run time sends a validate connection message to the remote end. Heartbeats are not supported on datagram connections.

If not defined, the default value is `1`.

Ice.ACM.Timeout

Synopsis

`Ice.ACM.Timeout=`*`num`*

Description

The value for *`num`* represents a timeout in seconds for Active Connection Management (ACM). The setting for `Ice.ACM.Close` determines when connections are automatically closed, and the setting for `Ice.ACM.Heartbeat` can be used to prevent connections from being closed prematurely. Use the corresponding `Ice.ACM.Client.Timeout` or `Ice.ACM.Server.Timeout` properties to override this setting in a client or server context, respectively.

If not defined, the default value is `60`.

Ice.ACM.Server.Close

Synopsis

`Ice.ACM.Server.Close=`*`num`*

Description

Overrides the value of `Ice.ACM.Close` in a server context, meaning for incoming connections.

Ice.ACM.Server.Heartbeat

Copyright 2016, ZeroC, Inc.

Synopsis

`Ice.ACM.Server.Heartbeat=`*`num`*

Description

Overrides the value of `Ice.ACM.Heartbeat` in a server context, meaning for incoming connections.

Ice.ACM.Server.Timeout

Synopsis

`Ice.ACM.Server.Timeout=`*`num`*

Description

Overrides the value of `Ice.ACM.Timeout` in a server context, meaning for incoming connections.

Copyright 2016, ZeroC, Inc.

Ice.Admin.*

On this page:

- Ice.Admin.AdapterProperty
- Ice.Admin.DelayCreation
- Ice.Admin.Enabled
- Ice.Admin.Facets
- Ice.Admin.InstanceName
- Ice.Admin.Logger.KeepLogs
- Ice.Admin.Logger.KeepTraces
- Ice.Admin.Logger.Properties
- Ice.Admin.ServerId

Ice.Admin.*AdapterProperty*

Synopsis

`Ice.Admin.`*`AdapterProperty`*`=`*`value`*

Description

The Ice run time creates and activates an administrative object adapter named `Ice.Admin` if the Administrative Facility is enabled, `Ice.Admin.Endpoints` is defined and one of the following are true:

- `Ice.Admin.DelayCreation` is not enabled
- `Ice.Admin.DelayCreation` is enabled and the application calls `getAdmin` on the communicator after communicator initialization
- the application calls `createAdmin` with a null `adminAdapter` parameter

This object adapter is created to host the admin object. Adapter properties can be used to configure the `Ice.Admin` object adapter.

Note that enabling the `Ice.Admin` object adapter is a security risk because a hostile client could use the administrative object to shut down the process. As a result, the endpoints for this object adapter should be carefully defined so that only trusted clients are allowed to use it.

Ice.Admin.DelayCreation

Synopsis

`Ice.Admin.DelayCreation=`*`num`*

Description

If *`num`* is a value greater than zero, the Ice run time delays the creation of the `Ice.Admin` administrative object adapter until `getAdmin` is invoked on the communicator. If not specified, the default value is zero, meaning the `Ice.Admin` object adapter is created immediately after all plug-ins are initialized, provided `Ice.Admin.Endpoints` is defined.

Ice.Admin.Enabled

Synopsis

`Ice.Admin.Enabled=`*`num`*

Description

If *`num`* is a value greater than zero, the Administrative Facility is enabled. If *`num`* is a zero or a negative value, the Administrative Facility is disabled. If this property is not set at all, the Administrative Facility is enabled when Ice.Admin.Endpoints is defined and not empty, and is disabled otherwise.

Ice.Admin.Facets

Copyright 2016, ZeroC, Inc.

Synopsis

`Ice.Admin.Facets=`*`name`* `[`*`name ...`*`]`

Description

Specifies the facets enabled by the administrative object, allowing you to filter the facets that the administrative object enables by default. Facet names are delimited by commas or white space. A facet name that contains white space must be enclosed in single or double quotes. If not specified, all facets are enabled. While the Ice run time creates only the built-in facets (such as Process and Properties) that are enabled, you can create administrative facets without checking the value of this property. Ice ensures that only enabled administrative facets are available to clients.

Ice.Admin.InstanceName

Synopsis

`Ice.Admin.InstanceName=`*`name`*

Description

Specifies an identity category for the administrative object, when this object is created during communicator initialization or by a call to `getAdmin` on the communicator. If defined, the identity of the object becomes *`name`*`/admin`. If not specified, the default identity category is a UUID.

Ice.Admin.Logger.KeepLogs

Synopsis

`Ice.Admin.Logger.KeepLogs=`*`num`*

Description

The `Logger` admin facet, when enabled, caches up the *num* most recent log messages with a type other than `Ice::TraceMessage`. When *num* is 0 or less than 0, the Logger facet does not cache any of these log messages. The default value for *num* is 100.

Ice.Admin.Logger.KeepTraces

Synopsis

`Ice.Admin.Logger.KeepTraces=`*`num`*

Description

The `Logger` admin facet, when enabled, caches up the *num* most recent log messages with type `Ice::TraceMessage`. When *num* is 0 or less than 0, the Logger facet does not cache any of these trace messages. The default value for *num* is 100.

Ice.Admin.Logger.Properties

Synopsis

`Ice.Admin.Logger.Properties=`*`propertyList`*

Description

The `Logger` admin facet, when enabled, creates its own communicator to send log messages to attached remote loggers. Without this sub-communicator, sending log messages to remote loggers could trigger more local logging, which in turn would generate more logs sent to remote loggers: a single genuine log could trigger an infinite number of log messages.

The properties of this sub-communicator are a few properties of the application's communicator (`Ice.Default.Locator`, `Ice.Plugin.IceSSL` and all `IceSSL` properties), plus the properties (if any) specified by *propertyList*. *propertyList* is a sequence of strings, that Ice reads

Copyright 2016, ZeroC, Inc.

using `getPropertyAsList`. Each of these strings uses the syntax *PropertyName=PropertyValue* to set a property. For example, you could turn on protocol tracing on the `Logger` facet's sub-communicator with:

```
Ice.Admin.Logger.Properties=Ice.Trace.Protocol=1
```

Ice.Admin.ServerId

Synopsis

`Ice.Admin.ServerId=`*`id`*

Description

Specifies an identifier that uniquely identifies the process when the Ice run time registers the Process facet of its admin object with the locator registry.

Copyright 2016, ZeroC, Inc.

Ice.Config

Ice.Config

Synopsis

```
--Ice.Config
--Ice.Config=1
--Ice.Config=config_file[,config_file,...]
```

Description

This property must be set from the command line with one of the options `--Ice.Config`, `--Ice.Config=1`, or `--Ice.Config=`*`config_file`*.

If the `Ice.Config` property is empty or set to 1, the Ice run time examines the contents of the `ICE_CONFIG` environment variable to retrieve the path names of one or more configuration files. Otherwise, `Ice.Config` must be set to the path names of one or more configuration files, separated by commas. (Path names can be relative or absolute.) Further property values are read from the configuration files thus specified.

In Java, Ice first attempts to open a configuration file as a class loader resource. If that attempt fails, Ice opens the configuration file in the local file system.

Configuration files use a simple syntax consisting of *name=value* pairs with support for comments and escaping.

See Also

- Using Configuration Files
- Configuration File Syntax
- Alternate Property Stores

Copyright 2016, ZeroC, Inc.

Ice.Default.*

On this page:

- Ice.Default.CollocationOptimized
- Ice.Default.EncodingVersion
- Ice.Default.EndpointSelection
- Ice.Default.Host
- Ice.Default.InvocationTimeout
- Ice.Default.Locator
- Ice.Default.LocatorCacheTimeout
- Ice.Default.Package
- Ice.Default.PreferSecure
- Ice.Default.Protocol
- Ice.Default.Router
- Ice.Default.SlicedFormat
- Ice.Default.SourceAddress
- Ice.Default.Timeout

Ice.Default.CollocationOptimized

Synopsis

`Ice.Default.CollocationOptimized=`*`num`*

Description

Specifies whether proxy invocations use collocation optimization by default. When enabled, proxy invocations on a collocated servant (i.e., a servant whose object adapter was created by the same communicator as the proxy) are made more efficiently by avoiding the network stack.

If not specified, the default value is 1. Set the property to 0 to disable collocation optimization by default.

Ice.Default.EncodingVersion

Synopsis

`Ice.Default.EncodingVersion=`*`ver`*

Description

If this property is not defined, Ice 3.5 uses encoding version 1.1 by default. To force the Ice run time to use encoding version 1.0 as the default instead, set this property to `1.0`:

```
Ice.Default.EncodingVersion=1.0
```

Ice.Default.EndpointSelection

Synopsis

`Ice.Default.EndpointSelection=`*`policy`*

Description

This property controls the default endpoint selection policy for proxies with multiple endpoints. Permissible values are `Ordered` and `Random`. The default value of this property is `Random`.

Ice.Default.Host

Synopsis

`Ice.Default.Host=`*`host`*

Copyright 2016, ZeroC, Inc.

Description

If an endpoint does not specify a host name (i.e., omits the `-h` *`host`* option), the *`host`* value from this property is used instead. The property has no default value.

Ice.Default.InvocationTimeout

Synopsis

`Ice.Default.InvocationTimeout=`*`num`*

Description

Specifies the default invocation timeout in milliseconds to use for all proxies. If not defined, the default timeout is `-1`, which means an invocation never times out.

If set to `-2`, invocation timeouts are disabled and the Ice run time behaves like Ice versions < 3.6: it uses connection timeouts (defined on the endpoints) to wait for the response of the invocation. This setting is provided only for backward compatibility and might be deprecated in a future Ice major release.

Ice.Default.Locator

Synopsis

`Ice.Default.Locator=`*`locator`*

Description

Specifies a default locator for all proxies and object adapters. The value is a stringified proxy for the IceGrid locator object. The default locator can be overridden on a proxy using the `ice_locator` proxy method. The default value is no locator.

The default identity of the IceGrid locator object is `IceGrid/Locator`, but this identity is influenced by the `IceGrid.InstanceName` property. The locator object is available on the IceGrid client endpoints. For example, suppose `IceGrid.Registry.Client.Endpoints` is set as follows:

```
IceGrid.Registry.Client.Endpoints=tcp -p 12000 -h localhost
```

In this case, the stringified proxy for the IceGrid locator is:

```
Ice.Default.Locator=IceGrid/Locator:tcp -p 12000 -h localhost
```

As a proxy property, you can configure additional aspects of the proxy using properties.

Ice.Default.LocatorCacheTimeout

Synopsis

`Ice.Default.LocatorCacheTimeout=`*`num`*

Description

Specifies the default locator cache timeout for indirect proxies. If *`num`* is set to a value larger than 0, locator cache entries older than *`num`* seconds are ignored. If set to 0, the locator cache is not used. If set to `-1`, locator cache entries do not expire.

Copyright 2016, ZeroC, Inc.

Once a cache entry has expired, the Ice run time performs a new locate request to refresh the cache before sending the next invocation; therefore, the invocation is delayed until the run time has refreshed the entry. If you set `Ice.BackgroundLocatorCacheUpdates` to a non-0value, the lookup to refresh the cache is still performed but happens in the background; this avoids the delay for the first invocation that follows expiry of a cache entry.

Ice.Default.Package

Synopsis

`Ice.Default.Package=`*`package`*

Description

Ice for Java allows you to customize the packaging of generated code. If you use this feature, the Ice run time requires additional configuration in order to successfully unmarshal exceptions and concrete class types. This property specifies a default package to use if other attempts by the Ice run time to dynamically load a generated class have failed. Also see `Ice.Package.`*`module`*.

Ice.Default.PreferSecure

Synopsis

`Ice.Default.PreferSecure=`*`num`*

Description

Specifies whether secure endpoints are given precedence in proxies by default. The default value of *`num`* is 0, meaning that insecure endpoints are given preference.

Setting this property to a non-0 value is the equivalent of invoking the proxy method `ice_preferSecure(true)` on proxies created by the Ice run time, such as those returned by `stringToProxy` or received as the result of an invocation. Proxies created by methods such as `ice_oneway` inherit the setting of the original proxy. If you want to force all proxies to use only secure endpoints, use `Ice.Override.Secure` instead.

See Configuring Secure Proxies for a discussion of secure proxies.

Ice.Default.Protocol

Synopsis

`Ice.Default.Protocol=`*`protocol`*

Description

Sets the protocol that is being used if an endpoint uses default as the protocol specification. The default value is `tcp`.

Ice.Default.Router

Synopsis

`Ice.Default.Router=`*`router`*

Description

Specifies the default router for all proxies. The value is a stringified proxy for the Glacier2 router control interface. The default router can be overridden on a proxy using the `ice_router` proxy method. The default value is no router.

As a proxy property, you can configure additional aspects of the proxy using properties.

Ice.Default.SlicedFormat

Copyright 2016, ZeroC, Inc.

Synopsis

`Ice.Default.SlicedFormat=`*`num`*

Description

Specifies the encoding format of Slice classes and exceptions. The default value of *`num`* is 0, meaning that the encoding uses the compact format. Set this property to a non-0 value to use the sliced format by default. This setting is only relevant when using version 1.1 of the Ice encoding.

Note that you can also specify whether certain operations use the sliced format by annotating their definitions with metadata.

Ice.Default.SourceAddress

Synopsis

`Ice.Default.SourceAddress=`*`addr`*

Description

If specified, outgoing socket connections are bound to the network interface associated with the given address *`addr`*, which must be a numeric IP address. Proxy endpoints can override this setting with the `--sourceAddress` option. If no source address is configured, the Ice run time uses the operating system's default behavior for binding an outgoing socket connection.

This feature is not supported on WinRT.

Ice.Default.Timeout

Synopsis

`Ice.Default.Timeout=`*`num`*

Description

Specifies the default timeout in milliseconds to use for an endpoint if no timeout is configured using the `-t` option. Use the value `-1` to configure an infinite default timeout. If not defined, the default timeout is `60000`.

Copyright 2016, ZeroC, Inc.

Ice.InitPlugins

Ice.InitPlugins

Synopsis

`Ice.InitPlugins=`*num*

Description

If *num* is a value greater than zero, the Ice run time automatically initializes the plug-ins it has loaded. The order in which plug-ins are loaded and initialized is determined by `Ice.PluginLoadOrder`. An application may need to set this property to zero in order to interact directly with a plug-in after it has been loaded but before it is initialized. In this case, the application must invoke `initializePlugins` on the plug-in manager to complete the initialization process. If not defined, the default value is 1.

When the admin Object is created, its `Process` facet is registered with the configured Locator registry if `Ice.Admin.ServerId` is set. This remote registration call may depend on plug-ins, such as the IceSSL plug-in. As a result, if you delay the initialization of the plug-ins by setting `Ice.InitPlugins` to 0, you may also want to delay the creation of the admin Object by setting `Ice.Admin.DelayCreation` to 1.

Copyright 2016, ZeroC, Inc.

Ice.IPv4

Ice.IPv4

Synopsis

`Ice.IPv4=`*`num`*

Description

Specifies whether Ice uses IPv4. If *`num`* is a value greater than zero, IPv4 is enabled. If not specified, the default value is 1.

Copyright 2016, ZeroC, Inc.

Ice.IPv6

Ice.IPv6

Synopsis

`Ice.IPv6=`*`num`*

Description

Specifies whether Ice uses IPv6. If *`num`* is a value greater than zero, IPv6 is enabled. If not specified, the default value is 1 if the system supports the creation of IPv6 sockets, and 0 otherwise.

Platform Notes

Java

Java's default network stack always accepts both IPv4 and IPv6 connections regardless of the settings of `Ice.IPv6`. You can configure the Java run time to use only IPv4 by starting your application with the following JVM option:

Java

```
java -Djava.net.preferIPv4Stack=true ...
```


Copyright 2016, ZeroC, Inc.

Ice.Override.*

On this page:

- Ice.Override.CloseTimeout
- Ice.Override.Compress
- Ice.Override.ConnectTimeout
- Ice.Override.Secure
- Ice.Override.Timeout

Ice.Override.CloseTimeout

Synopsis

`Ice.Override.CloseTimeout=`*`num`*

Description

This property overrides timeout settings used to close connections. *`num`* is the timeout value in milliseconds, or `-1` for no timeout. If this property is not defined, then `Ice.Override.Timeout` is used. If `Ice.Override.Timeout` is not defined, the endpoint timeout is used.

Ice.Override.Compress

Synopsis

`Ice.Override.Compress=`*`num`*

Description

If set, this property overrides compression settings in all proxies. If *`num`* is set to a value larger than zero, compression is enabled. If zero, compression is disabled.

The setting of this property is ignored in the server role.

Note that, if a client sets `Ice.Override.Compress=1` and sends a compressed request to a server that does not support compression, the server will close the connection and the client will receive `ConnectionLostException`.

If a client does not support compression and `Ice.Override.Compress=1`, the setting is ignored and a warning message is printed on `std err`.

Regardless of the setting of this property, requests smaller than 100 bytes are never compressed.

Ice.Override.ConnectTimeout

Synopsis

`Ice.Override.ConnectTimeout=`*`num`*

Description

This property overrides timeout settings used to establish connections. *`num`* is the timeout value in milliseconds, or `-1` for no timeout. If this property is not defined, then `Ice.Override.Timeout` is used. If `Ice.Override.Timeout` is not defined, the endpoint timeout is used.

Ice.Override.Secure

Synopsis

`Ice.Override.Secure=`*`num`*

Description

Copyright 2016, ZeroC, Inc.

If set to a value larger than zero, this property overrides security settings in all proxies by allowing only secure endpoints. Defining this property is equivalent to invoking the `ice_secure(true)` proxy method on every proxy. If you wish to give priority to secure endpoints without precluding the use of non-secure endpoints, use `Ice.Default.PreferSecure`. Refer to Configuring Secure Proxies for more information on secure proxies.

Ice.Override.Timeout

Synopsis

`Ice.Override.Timeout=`*`num`*

Description

If set, this property overrides timeout settings in all endpoints. *`num`* is the timeout value in milliseconds, or `-1` for no timeout.

Copyright 2016, ZeroC, Inc.

Ice.Plugin.*

On this page:

- Ice.Plugin.name
- Ice.Plugin.name.clr
- Ice.Plugin.name.cpp
- Ice.Plugin.name.java

Ice.Plugin.*name*

Synopsis

`Ice.Plugin.`*`name`*`=`*`entry_point`* `[`*`args`*`]`

Description

Defines a plug-in to be installed during communicator initialization. The format of *`entry_point`* varies by Ice implementation language, therefore this property cannot be defined in a configuration file that is shared by programs in different languages. Ice provides an alternate syntax that facilitates such sharing:

- `Ice.Plugin.`*`name`*`.cpp` for C++
- `Ice.Plugin.`*`name`*`.java` for Java
- `Ice.Plugin.`*`name`*`.clr` for the .NET Common Language Runtime

Refer to the relevant property for your language mapping for details on the entry point syntax.

Ice.Plugin.*name*.clr

Synopsis

`Ice.Plugin.`*`name`*`.clr=`*`assembly`*`:`*`class`* `[`*`args`*`]`

Description

Defines a .NET plug-in to be installed during communicator initialization. The *`assembly`* component can be a partially or fully qualified assembly name, such as `myplugin,Version=0.0.0.0,Culture=neutral`, or an assembly DLL name such as `myplugin.dll` that may optionally include a leading relative or absolute path name.

The locations that are searched for the assembly varies depending on how you define the *`assembly`* component:

Value for *`assembly`*	Example	Semantics
Fully-qualified assembly name (strong-named assembly)	`myplugin,Version=...,Culture=neutral,publicKeyToken=...`	1. Checks assemblies that have already been loaded 2. Searches the Global Assembly Cache (GAC) 3. Searches the directory containing the `iceboxnet` executable
Partially-qualified assembly name	`myplugin`	1. Checks assemblies that have already been loaded 2. Searches the directory containing the `iceboxnet` executable

Copyright 2016, ZeroC, Inc.

Relative path name	`plugins\MyPlugin.dll`	Path name is relative to the application's current working directory. Be sure to include the `.dll` extension in the path name.
Absolute path name	`C:\plugins\MyPlugin.dll`	Assembly must reside at the specified path name. Be sure to include the `.dll` extension in the path name.

See MSDN for more information on how the CLR locates assemblies.

The specified *`class`* must implement the `Ice.PluginFactory` interface. Any arguments that follow the class name are passed to the factory's `create` method. For example:

```
Ice.Plugin.MyPlugin.clr=MyFactory,Version=1.2.3.4:MyFactory arg1 arg2
```

Whitespace separates the arguments, and any arguments that contain whitespace must be enclosed in quotes.

If you specify a relative path name in the entry point, the assembly is located relative to the program's current working directory:

```
Ice.Plugin.MyPlugin.clr=..\MyFactory.dll:MyFactory arg1 arg2
```

Enclose the assembly's path name in quotes if it contains spaces:

```
Ice.Plugin.MyPlugin.clr="C:\Program
Files\MyPlugin\MyFactory.dll:MyFactory" arg1 arg2
```

Ice.Plugin.*name*.cpp

Synopsis

`Ice.Plugin.`*`name`*`.cpp=`*`path`*`[,`*`version`*`]:`*`function`* `[`*`args`*`]`

Description

Defines a C++ plug-in to be installed during communicator initialization. The *`path`* and optional *`version`* components are used to construct the path name of a DLL or shared library. If no version is supplied, the Ice version is used. The *`function`* component is the name of a function with C linkage. For example, the entry point `MyPlugin,36:create` would imply a shared library name of `libMyPlugin.so.36` on Unix and `MyPlugin36.dll` on Windows. Furthermore, if Ice is built on Windows with debugging, a `d` is automatically appended to the version (for example, `MyPlugin36d.dll`).

The function must be declared with external linkage and have the following signature:

Copyright 2016, ZeroC, Inc.

C++

```
Ice::Plugin* function(const Ice::CommunicatorPtr& communicator,
                      const std::string& name,
                      const Ice::StringSeq& args);
```

Note that the function must return a pointer and not a smart pointer.

Any arguments that follow the entry point are passed to the entry point function. For example:

```
Ice.Plugin.MyPlugin.cpp=MyFactory,36:create arg1 arg2
```

Whitespace separates the arguments, and any arguments that contain whitespace must be enclosed in quotes.

The *path* component may optionally contain a relative or absolute path name, indicated by the presence of a path separator (`/` or `\`). In this case, the last component of the path is used to construct the version-specific name of the shared library or DLL. Consider this example:

```
Ice.Plugin.MyPlugin.cpp=./MyFactory,36:create arg1 arg2
```

The use of a relative path means the Ice run time will look in the current working directory for `libMyPlugin.so.36` on Unix or `MyPlugin36.dll` on Windows.

If the *path* component contains spaces, the entire entry point must be enclosed in quotes:

```
Ice.Plugin.MyPlugin.cpp="C:\Program Files\MyPlugin\MyFactory,36:create" arg1 arg2
```

If the *path* component does not include a leading path name, Ice delegates to the operating system to locate the shared library or DLL, which typically means that the plug-in can reside in any of the directories in your shared library or DLL search path.

When the plug-in is packaged in a static library and linked into the application through `Ice::registerPluginFactory`, the entry point (*path*[,*version*]:*function)* component of this property is ignored. The `args`, if any, are preserved, and are given to the registered plug-in factory function when the plug-in is created.

Ice.Plugin.*name*.java

Synopsis

`Ice.Plugin.`*`name`*`.java=[`*`path`*`:]`*`class`* `[`*`args`*`]`

Description

Defines a Java plug-in to be installed during communicator initialization. The specified class must implement the `Ice.PluginFactory` inter face. Any arguments that follow the class name are passed to the `create` method. For example:

Copyright 2016, ZeroC, Inc.

```
Ice.Plugin.MyPlugin.java=MyFactory arg1 arg2
```

Whitespace separates the arguments, and any arguments that contain whitespace must be enclosed in quotes.

If *path* is specified, it may be the path name of a JAR file or class directory, as shown below:

```
Ice.Plugin.MyPlugin.java=MyFactory.jar:MyFactory
Ice.Plugin.MyOtherPlugin.java=/classes:MyOtherFactory
```

If *path* contains spaces, it must be enclosed in quotes:

```
Ice.Plugin.MyPlugin.java="factory classes.jar":MyFactory
```

If *class* is specified without a path, Ice attempts to load the class using class loaders in a well-defined order.

Copyright 2016, ZeroC, Inc.

Ice.PluginLoadOrder

Ice.PluginLoadOrder

Synopsis

`Ice.PluginLoadOrder=`*`names`*

Description

Determines the order in which plug-ins are loaded (loaded is a synonym for created in this context). The Ice run time loads the plug-ins in the order they appear in *`names`*, where each plug-in name is separated by a comma or white space. Any plug-ins not mentioned in *`names`* are loaded afterward, in an undefined order.

C++ plug-ins for which factories were registered through `Ice::registerPluginFactory` with the last parameter set to `true` are always loaded first, in registration order, and are not affected by this property.

Copyright 2016, ZeroC, Inc.

Ice.PreferIPv6Address

Ice.PreferIPv6Address

Synopsis

`Ice.PreferIPv6Address=`*`num`*

Description

If both IPv4 and IPv6 are enabled (the default), specifies whether Ice prefers IPv6 addresses over IPv4 addresses when resolving hostnames. If *`num`* is a value greater than zero, IPv6 addresses are preferred. If not specified, the default value is 0.

Copyright 2016, ZeroC, Inc.

Ice.TCP.*

On this page:

- Ice.TCP.Backlog
- Ice.TCP.RcvSize
- Ice.TCP.SndSize

Ice.TCP.Backlog

Synopsis

`Ice.TCP.Backlog=`*`num`*

Description

Specifies the size of the listen queue for each TCP or SSL server endpoint. If not defined, the default value for C++ programs uses the value of `SOMAXCONN` if present, or `511` otherwise. In Java and .NET, the default value is 511.

Ice.TCP.RcvSize

Synopsis

`Ice.TCP.RcvSize=`*`num`*

Description

This property sets the TCP receive buffer size to the specified value in bytes. The default value depends on the configuration of the local TCP stack. (A common default values is 65535 bytes.)

The OS may impose lower and upper limits on the receive buffer size or otherwise adjust the buffer size. If a limit is requested that is lower than the OS-imposed minimum, the value is silently adjusted to the OS-imposed minimum. If a limit is requested that is larger than the OS-imposed maximum, the value is adjusted to the OS-imposed maximum; in addition, Ice logs a warning showing the requested size and the adjusted size.

Ice.TCP.SndSize

Synopsis

`Ice.TCP.SndSize=`*`num`*

Description

This property sets the TCP send buffer size to the specified value in bytes. The default value depends on the configuration of the local TCP stack. (A common default values is 65535 bytes.)

The OS may impose lower and upper limits on the send buffer size or otherwise adjust the buffer size. If a limit is requested that is lower than the OS-imposed minimum, the value is silently adjusted to the OS-imposed minimum. If a limit is requested that is larger than the OS-imposed maximum, the value is adjusted to the OS-imposed maximum; in addition, Ice logs a warning showing the requested size and the adjusted size.

Copyright 2016, ZeroC, Inc.

Ice.ThreadPool.*

A communicator creates two thread pools: the client thread pool dispatches AMI callbacks and incoming requests on bidirectional connections, and the server thread pool dispatches requests to object adapters.

This page describes configuration properties for the client and server thread pools. These thread pools are named `Client` and `Server`, respectively. In the property descriptions below, replace *name* with `Client` or `Server`.

On this page:

- Ice.ThreadPool.name.Serialize
- Ice.ThreadPool.name.Size
- Ice.ThreadPool.name.SizeMax
- Ice.ThreadPool.name.SizeWarn
- Ice.ThreadPool.name.StackSize
- Ice.ThreadPool.name.ThreadIdleTime
- Ice.ThreadPool.name.ThreadPriority

Ice.ThreadPool.*name*.Serialize

Synopsis

`Ice.ThreadPool.`*`name`*`.Serialize=`*`num`*

Description

If *num* is a value greater than 0, the `Client` or `Server` thread pool serializes all messages from each connection. It is not necessary to enable this feature in a thread pool whose maximum size is 1 thread. In a multi-threaded pool, enabling serialization allows requests from different connections to be dispatched concurrently while preserving the order of messages on each connection. Note that serialization has a signficant impact on latency and throughput. If not defined, the default value is 0.

Ice.ThreadPool.*name*.Size

Synopsis

`Ice.ThreadPool.`*`name`*`.Size=`*`num`*

Description

Thread pools in Ice can grow and shrink dynamically, based on an average load factor. A thread pool always has at least 1 thread and may grow as load increases up to the maximum size specified by `Ice.ThreadPool.`*`name`*`.SizeMax`. If `SizeMax` is not specified, Ice uses the value of *num* as the pool's maximum size. The `Client` or `Server` thread pool is initialized with *num* active threads, but the pool may shrink to only 1 thread during idle periods as determined by `Ice.ThreadPool.`*`name`*`.ThreadIdleTime`.

If not specified, the default value is 1 for both properties.

An object adapter can also be configured with its own thread pool.

Note that multiple threads for the client thread pool are only required for nested AMI invocations, or to allow multiple AMI callbacks to be dispatched concurrently.

To monitor the thread pool activities of the Ice run time, enable the `Ice.Trace.ThreadPool` property.

Ice.ThreadPool.*name*.SizeMax

Synopsis

`Ice.ThreadPool.`*`name`*`.SizeMax=`*`num`*

Description

Copyright 2016, ZeroC, Inc.

num is the maximum number of threads for the `Client` or `Server` thread pool. Refer to the `Ice.ThreadPool.`*name*`.Size` property for more information on configuring the size of a thread pool.

The default value for `SizeMax` is the value of `Size`, meaning the thread pool can never grow larger than its initial size.

To monitor the thread pool activities of the Ice run time, enable the `Ice.Trace.ThreadPool` property.

Ice.ThreadPool.*name*.SizeWarn

Synopsis

`Ice.ThreadPool.`*name*`.SizeWarn=`*num*

Description

Whenever *num* threads are active in the `Client` or `Server` thread pool, a "low on threads" warning is printed. The default value is 0, which disables the warning.

To monitor the thread pool activities of the Ice run time, enable the `Ice.Trace.ThreadPool` property.

Ice.ThreadPool.*name*.StackSize

Synopsis

`Ice.ThreadPool.`*name*`.StackSize=`*num*

Description

num is the stack size (in bytes) of threads in the `Client` or `Server` thread pool. The default value is 0, meaning the operating system's default is used.

Ice.ThreadPool.*name*.ThreadIdleTime

Synopsis

`Ice.ThreadPool.`*name*`.ThreadIdleTime=`*num*

Description

Ice can automatically reap idle threads in the `Client` or `Server` thread pool to conserve resources. This property specifies the number of seconds a thread must be idle before it is reaped. If not specified, the default value is 60 seconds.

To disable the reaping of idle threads, set `ThreadIdleTime` to 0. In this situation, the thread pool is initialized with `Ice.ThreadPool.`*name*`.Size` active threads and may grow to contain `Ice.ThreadPool.`*name*`.SizeMax` active threads, but the size of the pool never decreases.

To monitor the thread pool activities of the Ice run time, enable the `Ice.Trace.ThreadPool` property.

Ice.ThreadPool.*name*.ThreadPriority

Synopsis

`Ice.ThreadPool.`*name*`.ThreadPriority=`*num*

Description

num specifies a thread priority for the threads in the `Client` or `Server` thread pool. Leaving this property unset causes the run time to create threads with the default priority specified by `Ice.ThreadPriority`.

Copyright 2016, ZeroC, Inc.

This property is unset by default.

You can also override the default priority for a specific object adapter using *adapter*`.ThreadPool.ThreadPriority`.

Copyright 2016, ZeroC, Inc.

Ice.ThreadPriority

Ice.ThreadPriority

Synopsis

`Ice.ThreadPriority=`*`num`*

Description

`num` specifies a thread priority. Threads created by the Ice run time are created with the specified priority by default. Leaving this property unset causes the run time to create threads with the system default priority. This property is unset by default.

You can separately override the default priorities for the client and server thread pools using `Ice.ThreadPool.`*`name`*`.ThreadPriority` a s well as for a specific object adapter using *`adapter`*`.ThreadPool.ThreadPriority`.

Copyright 2016, ZeroC, Inc.

Ice.Trace.*

On this page:

- Ice.Trace.Admin.Logger
- Ice.Trace.Admin.Properties
- Ice.Trace.Locator
- Ice.Trace.Network
- Ice.Trace.Protocol
- Ice.Trace.Retry
- Ice.Trace.Slicing
- Ice.Trace.ThreadPool

Ice.Trace.Admin.Logger

Synopsis

`Ice.Trace.Admin.Logger=`*`num`*

Description

Controls the trace level for the Logger administrative facet.

0	No trace (default).
1	Trace when a remote logger is attached or detached.
2	Like 1, but also trace the sending of log messages to remote loggers.

Ice.Trace.Admin.Properties

Synopsis

`Ice.Trace.Admin.Properties=`*`num`*

Description

Controls the trace level for property updates made via the `Properties` facet:

0	No property trace (default).
1	Trace property addition, modification, and removal.

Ice.Trace.Locator

Synopsis

`Ice.Trace.Locator=`*`num`*

Description

The Ice run time makes locator requests to resolve the endpoints of object adapters and well-known objects. Requests on the locator registry are used to update object adapter endpoints and set the server process proxy. This property controls the trace level for the Ice run time's interactions with the locator:

0	No locator trace (default).
1	Trace Ice locator and locator registry requests.
2	Like 1, but also trace the removal of endpoints from the cache.

Copyright 2016, ZeroC, Inc.

Ice.Trace.Network

Synopsis

`Ice.Trace.Network=`*`num`*

Description

Controls the trace level for low-level network activities such as connection establishment and read/write operations:

0	No network trace (default).
1	Trace successful connection establishment and closure.
2	Like 1, but also trace attempts to bind, connect, and disconnect sockets.
3	Like 2, but also trace data transfer, the published endpoints for an object adapter, and the current list of local addresses for an endpoint that uses the wildcard address.

Ice.Trace.Protocol

Synopsis

`Ice.Trace.Protocol=`*`num`*

Description

Controls the trace level for Ice protocol mesages:

0	No protocol trace (default).
1	Trace Ice protocol messages.

Ice.Trace.Retry

Synopsis

`Ice.Trace.Retry=`*`num`*

Description

Ice supports automatic retries in case of a request failure. This property controls the trace level for retry attempts:

0	No request retry trace (default).
1	Trace Ice operation call retries.
2	Also trace Ice endpoint usage.

Ice.Trace.Slicing

Synopsis

`Ice.Trace.Slicing=`*`num`*

Description

The Ice data encoding for exceptions and classes enables a receiver to slice an unknown exception or class type to a known type. This property controls the trace level for slicing activities:

Copyright 2016, ZeroC, Inc.

0	No trace of slicing activity (default).
1	Trace all exception and class types that are unknown to the receiver and therefore sliced.

Ice.Trace.ThreadPool

Synopsis

`Ice.Trace.ThreadPool=`*`num`*

Description

Controls the trace level for the Ice thread pool:

0	No trace of thread pool activity (default).
1	Trace the creation, growing, and shrinking of thread pools.

Copyright 2016, ZeroC, Inc.

Ice.UDP.*

On this page:

- Ice.UDP.RcvSize
- Ice.UDP.SndSize

Ice.UDP.RcvSize

Synopsis

`Ice.UDP.RcvSize=`*`num`*

Description

This property sets the UDP receive buffer size to the specified value in bytes. Ice messages larger than *`num`* `- 28` bytes cause a `DatagramLimitException`. The default value depends on the configuration of the local UDP stack. (Common default values are 65535 and 8192 bytes.)

The OS may impose lower and upper limits on the receive buffer size or otherwise adjust the buffer size. If a limit is requested that is lower than the OS-imposed minimum, the value is silently adjusted to the OS-imposed minimum. If a limit is requested that is larger than the OS-imposed maximum, the value is adjusted to the OS-imposed maximum; in addition, Ice logs a warning showing the requested size and the adjusted size.

Values less than 28 are ignored.

Note that, on many operating systems, it is possible to set a buffer size greater than 65535. Such settings do not change the hard limit of 65507 bytes for the payload of a UDP packet, but merely affect how much data can be buffered by the kernel.

Settings less than 65535 limit the size of Ice datagrams as well as adjust the kernel buffer sizes.

Ice.UDP.SndSize

Synopsis

`Ice.UDP.SndSize=`*`num`*

Description

This property sets the UDP send buffer size to the specified value in bytes. Ice messages larger than *`num`* `- 28` bytes cause a `DatagramLimitException`. The default value depends on the configuration of the local UDP stack. (Common default values are 65535 and 8192 bytes.)

The OS may impose lower and upper limits on the send buffer size or otherwise adjust the buffer size. If a limit is requested that is lower than the OS-imposed minimum, the value is silently adjusted to the OS-imposed minimum. If a limit is requested that is larger than the OS-imposed maximum, the value is adjusted to the OS-imposed maximum; in addition, Ice logs a warning showing the requested size and the adjusted size.

Values less than 28 are ignored.

Note that, on many operating systems, it is possible to set a buffer size greater than 65535. Such settings do not change the hard limit of 65507 bytes for the payload of a UDP packet, but merely affect how much data can be buffered by the kernel.

Settings less than 65535 limit the size of Ice datagrams as well as adjust the kernel buffer sizes.

Copyright 2016, ZeroC, Inc.

Ice.Warn.*

On this page:

- Ice.Warn.AMICallback
- Ice.Warn.Connections
- Ice.Warn.Datagrams
- Ice.Warn.Dispatch
- Ice.Warn.Endpoints
- Ice.Warn.UnknownProperties
- Ice.Warn.UnusedProperties

Ice.Warn.AMICallback

Synopsis

`Ice.Warn.AMICallback=`*`num`*

Description

If *`num`* is set to a value larger than zero, warnings are printed if an AMI callback raises an exception. The default value is 1.

Ice.Warn.Connections

Synopsis

`Ice.Warn.Connections=`*`num`*

Description

If *`num`* is set to a value larger than zero, Ice applications print warning messages for certain exceptional conditions in connections. The default value is 0.

Ice.Warn.Datagrams

Synopsis

`Ice.Warn.Datagrams=`*`num`*

Description

If *`num`* is set to a value larger than zero, a server prints a warning message if it receives a datagram that exceeds the server's receive buffer size. (Note that this condition is not detected by all UDP implementations — some implementations silently drop received datagrams that are too large.) The default value is 0.

Ice.Warn.Dispatch

Synopsis

`Ice.Warn.Dispatch=`*`num`*

Description

If *`num`* is set to a value larger than zero, Ice applications print warning messages for certain exceptions that are raised while an incoming request is dispatched.

Warning levels:

0	No warnings.

Copyright 2016, ZeroC, Inc.

1	Print warnings for unexpected `Ice::LocalException`, `Ice::UserException`, C++ exceptions, and Java run-time exceptions (default).
2	Like 1, but also issue warnings for `Ice::ObjectNotExistException`, `Ice::FacetNotExistException`, and `Ice::OperationNotExistException`.

Ice.Warn.Endpoints

Synopsis

`Ice.Warn.Endpoints=`*`num`*

Description

If *`num`* is set to a value larger than zero, a warning is printed if a stringified proxy contains an endpoint that cannot be parsed. (For example, stringified proxies containing SSL endpoints cause this warning in versions of Ice that do not support SSL.) The default value is 1.

Ice.Warn.UnknownProperties

Synopsis

`Ice.Warn.UnknownProperties=`*`num`*

Description

If *`num`* is set to a value larger than zero, the Ice run time prints a warning about unknown properties for object adapters and proxies. The default value is 1.

Ice.Warn.UnusedProperties

Synopsis

`Ice.Warn.UnusedProperties=`*`num`*

Description

If *`num`* is set to a value larger than zero, the Ice run time prints a warning during communicator destruction about properties that were set but not read. This warning is useful for detecting mis-spelled properties, such as `Filesystem.MaxFilSize`. The default value is 0.

Copyright 2016, ZeroC, Inc.

IceBox.*

On this page:

- IceBox.InheritProperties
- IceBox.LoadOrder
- IceBox.PrintServicesReady
- IceBox.Service.name
- IceBox.UseSharedCommunicator.name

IceBox.InheritProperties

Synopsis

`IceBox.InheritProperties=`*`num`*

Description

If *`num`* is set to a value larger than zero, each service inherits the configuration properties of the IceBox server's communicator. If not defined, the default value is zero.

IceBox.LoadOrder

Synopsis

`IceBox.LoadOrder=`*`names`*

Description

Determines the order in which services are loaded. The service manager loads the services in the order they appear in *`names`*, where each service name is separated by a comma or white space. Any services not mentioned in *`names`* are loaded afterward, in an undefined order.

IceBox.PrintServicesReady

Synopsis

`IceBox.PrintServicesReady=`*`token`*

Description

If this property is set to a value greater than zero, the service manager prints "*`token`* ready" on standard output once initialization of all the services is complete. This is useful for scripts that need to wait until all services are ready to be used.

IceBox.Service.*name*

Synopsis

`IceBox.Service.`*`name=entry_point`* `[`*`args`*`]`

Description

Defines a service to be loaded during IceBox initialization. Any arguments that follow the entry point are examined; those matching the `--`*`name`*`.*=`*`value`* pattern are interpreted as property definitions and appear in the property set of the communicator that is passed to the service `start` method, and all remaining arguments are passed to the `start` method in the `args` parameter. Whitespace separates the arguments, and any arguments that contain whitespace must be enclosed in quotes.

Platform Notes

Copyright 2016, ZeroC, Inc.

C++

The value of *entry_point* has the following form:

path[,*version*]:*function*

The *path* and optional *version* components are used to construct the name of a DLL or shared library. If no version is supplied, the version is the empty string. The *function* component is the name of a function with extern C linkage. For example, the entry point `IceStormService,35:createIceStorm` implies a shared library name of `libIceStormService.so.35` on Unix and `IceStormService35.dll` on Windows. Furthermore, if IceBox is built on Windows with debugging, a `d` is automatically appended to the version (e.g., `IceStormService35d.dll`).

The function must be declared with extern C linkage and have the following signature:

C++

```
IceBox::Service* function(Ice::CommunicatorPtr c);
```

Note that the function must return a pointer and not a smart pointer. The Ice core deallocates the object when it unloads the library. The communicator instance passed to this function is the server's communicator, which is not the same as the communicator passed to the service's `start` method.

The *path* component may optionally contain a relative or absolute path name, indicated by the presence of a path separator (`/` or `\`). In this case, the last component of the path is used to construct the name of the shared library or DLL. Consider this example:

```
IceBox.Service.IceStorm=./IceStormService,35:createIceStorm
```

The use of a relative path means the Ice run time will look in the current working directory for `libIceStormService.so.35` on Unix or `IceStormService35.dll` on Windows.

If the *path* component contains spaces, the entire entry point must be enclosed in quotes:

```
IceBox.Service.IceStorm="C:\Program
Files\ZeroC\Ice-3.5b\bin\IceStormService,35:createIceStorm"
```

If the *path* component does not include a leading path name, Ice delegates to the operating system to locate the shared library or DLL, which typically means that the plug-in can reside in any of the directories in your shared library or DLL search path.

Java

The value of *entry_point* has the following form:

[*path*:]*class*

The *class* component must be the name of a class that implements the `IceBox.Service` interface and provides at least one of the constructors shown in the example below:

Copyright 2016, ZeroC, Inc.

Java

```
public class MyService implements IceBox.Service
{
    public MyService(Ice.Communicator serverCommunicator);
    public MyService();

    // ...
}
```

The constructor taking an `Ice.Communicator` argument is invoked if present, otherwise the default constructor is invoked.

If *`path`* is specified, it may be the path name of a JAR file or class directory, as shown below:

```
IceBox.Service.MyService=MyService.jar:MyServiceImpl
IceBox.Service.MyOtherService=/classes:MyOtherServiceImpl
```

If *`path`* contains spaces, it must be enclosed in quotes:

```
IceBox.Service.MyService="factory classes.jar":MyServiceImpl
```

IceBox uses a single class loader to load all services having the same value for *`path`*.

If *`class`* is specified without a path, IceBox attempts to load the class using class loaders in a well-defined order.

.NET

The value of *`entry_point`* has the form *`assembly`*`:`*`class`*. The *`assembly`* can be a partially or fully qualified assembly name, such as `myplugin,Version=0.0.0.0,Culture=neutral`, or an assembly DLL name such as `myplugin.dll`, and may optionally include a leading relative or absolute path name.

You *must* use a fully-qualified assembly name to load a service from an assembly in the Global Assembly Cache.

The specified class must implement the `IceBox.Service` interface and provide at least one of the constructors shown in the example below:

C#

```
public class MyService : IceBox.Service
{
    public MyService(Ice.Communicator serverCommunicator);
    public MyService();

    // ...
}
```


Copyright 2016, ZeroC, Inc.

The constructor taking an `Ice.Communicator` argument is invoked if present, otherwise the default constructor is invoked.

If you specify a relative path name in the entry point, the assembly is located relative to the program's current working directory:

```
IceBox.Service.MyService=..\MyService.dll:MyServiceImpl
```

Enclose the assembly's path name in quotes if it contains spaces:

```
IceBox.Service.MyService="C:\Program Files\MyService\MyService.dll:MyServiceImpl"
```

Finally, if the assembly uses a leading path name, be sure to include the `.dll` extension.

IceBox.UseSharedCommunicator.*name*

Synopsis

`IceBox.UseSharedCommunicator.`*`name`*`=`*`num`*

Description

If *`num`* is set to a value larger than zero, the service manager supplies the service *`name`* with a communicator that might be shared by other services. If the `IceBox.InheritProperties` property is also defined, the shared communicator inherits the properties of the IceBox server. If not defined, the default value is zero.

Copyright 2016, ZeroC, Inc.

IceBoxAdmin

IceBoxAdmin.ServiceManager.Proxy

Synopsis

`IceBoxAdmin.ServiceManager.Proxy=`*`proxy`*

Description

This property configures the proxy that is used by the `iceboxadmin` utility to locate the service manager.

Copyright 2016, ZeroC, Inc.

IceDiscovery.*

This page describes the properties supported by the IceDiscovery plug-in.

On this page:

- IceDiscovery.Address
- IceDiscovery.DomainId
- IceDiscovery.Interface
- IceDiscovery.Locator.AdapterProperty
- IceDiscovery.Lookup
- IceDiscovery.Multicast.AdapterProperty
- IceDiscovery.Port
- IceDiscovery.Reply.AdapterProperty
- IceDiscovery.RetryCount
- IceDiscovery.Timeout

IceDiscovery.Address

Synopsis

`IceDiscovery.Address=`*`addr`*

Description

Specifies the multicast IP address to use for sending or receiving multicast requests. If not defined, the default value depends on other property settings:

- If `Ice.PreferIPv6Address` is enabled or `Ice.IPv4` is disabled, IceDiscovery uses the IPv6 address `ff15::1`
- Otherwise IceDiscovery uses `239.255.0.1`

IceDiscovery.DomainId

Synopsis

`IceDiscovery.DomainId=`*`id`*

Description

Specifies the domain ID used to locate objects and object adapters. The IceDiscovery plug-in only responds to requests from clients with the same domain ID and ignores requests from clients with a different domain ID. If not defined, the default domain ID is an empty string.

IceDiscovery.Interface

Synopsis

`IceDiscovery.Interface=`*`intf`*

Description

Specifies the IP address of the interface to use for sending or receiving multicast requests. If not defined, the operating system will select a default interface to send and receive the UDP multicast datagrams.

IceDiscovery.Locator.*AdapterProperty*

Synopsis

`IceDiscovery.Locator.`*`AdapterProperty=value`*

Description

Copyright 2016, ZeroC, Inc.

IceDiscovery creates an object adapter named `IceDiscovery.Locator`.

You don't normally need to set other properties for this object adapter.

IceDiscovery.Lookup

Synopsis

`IceDiscovery.Lookup=`*`endpoint`*

Description

Specifies the multicast endpoint that a client uses to send discovery requests. If not defined, the endpoint is composed as follows:

```
udp -h addr -p port --interface intf
```

where *`addr`* is the value of `IceDiscovery.Address`, *`port`* is the value of `IceDiscovery.Port`, and *`intf`* is the value of `IceDiscovery.Interface`.

IceDiscovery.Multicast.*AdapterProperty*

Synopsis

`IceDiscovery.Multicast.`*`AdapterProperty=value`*

Description

IceDiscovery creates an object adapter named `IceDiscovery.Multicast` for receiving discovery requests from clients. If not otherwise defined by `IceDiscovery.Multicast.Endpoints`, the endpoint for this object adapter is composed as follows:

```
udp -h addr -p port --interface intf
```

where *`addr`* is the value of `IceDiscovery.Address`, *`port`* is the value of `IceDiscovery.Port`, and *`intf`* is the value of `IceDiscovery.Interface`.

You don't normally need to set other properties for this object adapter.

IceDiscovery.Port

Synopsis

`IceDiscovery.Port=`*`port`*

Description

Specifies the multicast port to use for sending or receiving multicast requests. If not set, the default value is `4061`.

IceDiscovery.Reply.*AdapterProperty*

Synopsis

`IceDiscovery.Reply.`*`AdapterProperty=value`*

Description

IceDiscovery creates an object adapter named `IceDiscovery.Reply` for receiving replies to multicast requests. If not otherwise defined by `IceDiscovery.Reply.Endpoints`, the endpoint for this object adapter is composed as follows:

```
udp --interface intf
```


Copyright 2016, ZeroC, Inc.

where *intf* is the value of `IceDiscovery.Interface`. A fixed port is not necessary for this endpoint.

You don't normally need to set other properties for this object adapter.

IceDiscovery.RetryCount

Synopsis

`IceDiscovery.RetryCount=`*num*

Description

Specifies the maximum number of times that the plug-in will retry sending UDP multicast requests before giving up. The `IceDiscovery.Timeout` property determines how long the plug-in waits for a reply before trying again. If not defined, the default retry count is 3.

IceDiscovery.Timeout

Synopsis

`IceDiscovery.Timeout=`*num*

Description

Specifies the time interval in milliseconds to wait for replies to UDP multicast requests. If no server replies during this time interval, the client will retry the request the number of times specified by `IceDiscovery.RetryCount`. If not defined, the default timeout is `300`.

Copyright 2016, ZeroC, Inc.

IceGrid.*

On this page:

- IceGrid.InstanceName
- IceGrid.Node.AdapterProperty
- IceGrid.Node.AllowEndpointsOverride
- IceGrid.Node.AllowRunningServersAsRoot
- IceGrid.Node.CollocateRegistry
- IceGrid.Node.Data
- IceGrid.Node.DisableOnFailure
- IceGrid.Node.Name
- IceGrid.Node.Output
- IceGrid.Node.PrintServersReady
- IceGrid.Node.ProcessorSocketCount
- IceGrid.Node.PropertiesOverride
- IceGrid.Node.RedirectErrToOut
- IceGrid.Node.Trace.Activator
- IceGrid.Node.Trace.Adapter
- IceGrid.Node.Trace.Patch
- IceGrid.Node.Trace.Replica
- IceGrid.Node.Trace.Server
- IceGrid.Node.UserAccountMapper
- IceGrid.Node.UserAccounts
- IceGrid.Node.WaitTime
- IceGrid.Registry.AdminCryptPasswords
- IceGrid.Registry.AdminPermissionsVerifier
- IceGrid.Registry.AdminSessionFilters
- IceGrid.Registry.AdminSessionManager.AdapterProperty
- IceGrid.Registry.AdminSSLPermissionsVerifier
- IceGrid.Registry.Client.AdapterProperty
- IceGrid.Registry.CryptPasswords
- IceGrid.Registry.Data
- IceGrid.Registry.DefaultTemplates
- IceGrid.Registry.Discovery.AdapterProperty
- IceGrid.Registry.Discovery.Address
- IceGrid.Registry.Discovery.Enabled
- IceGrid.Registry.Discovery.Interface
- IceGrid.Registry.Discovery.Port
- IceGrid.Registry.DynamicRegistration
- IceGrid.Registry.Internal.AdapterProperty
- IceGrid.Registry.NodeSessionTimeout
- IceGrid.Registry.PermissionsVerifier
- IceGrid.Registry.ReplicaName
- IceGrid.Registry.ReplicaSessionTimeout
- IceGrid.Registry.Server.AdapterProperty
- IceGrid.Registry.SessionFilters
- IceGrid.Registry.SessionManager.AdapterProperty
- IceGrid.Registry.SessionTimeout
- IceGrid.Registry.SSLPermissionsVerifier
- IceGrid.Registry.Trace.Adapter
- IceGrid.Registry.Trace.Application
- IceGrid.Registry.Trace.Locator
- IceGrid.Registry.Trace.Node
- IceGrid.Registry.Trace.Object
- IceGrid.Registry.Trace.Patch
- IceGrid.Registry.Trace.Replica
- IceGrid.Registry.Trace.Server
- IceGrid.Registry.Trace.Session
- IceGrid.Registry.UserAccounts

IceGrid.InstanceName

Synopsis

`IceGrid.InstanceName=`*`name`*

Copyright 2016, ZeroC, Inc.

Description

Specifies an alternate identity category for the well-known IceGrid objects. If defined, the identities of the IceGrid objects become:

```
name/AdminSessionManager
name/AdminSessionManager-replica
name/AdminSSLSessionManager
name/AdminSSLSessionManager-replica
name/NullPermissionsVerifier
name/NullSSLPermissionsVerifier
name/Locator
name/Query
name/Registry
name/Registry-replica
name/RegistryUserAccountMapper
name/RegistryUserAccountMapper-replica
name/SessionManager
name/SSLSessionManager
```

If not specified, the default identity category is `IceGrid`.

IceGrid.Node.*AdapterProperty*

Synopsis

`IceGrid.Node.`*`AdapterProperty`*`=`*`value`*

Description

An IceGrid node uses the adapter name `IceGrid.Node` for the object adapter that the registry contacts to communicate with the node. Therefore, adapter properties can be used to configure this adapter.

IceGrid.Node.AllowEndpointsOverride

Synopsis

`IceGrid.Node.AllowEndpointsOverride=`*`num`*

If *`num`* is set to a non-zero value, an IceGrid node permits servers to override previously set endpoints even if the server is active. Setting this property to a non-zero value is necessary if the servers managed by the node use the object adapter operation `refreshPublishedEndpoints`. The default value of *`num`* is zero.

IceGrid.Node.AllowRunningServersAsRoot

Synopsis

`IceGrid.Node.AllowRunningServersAsRoot=`*`num`*

If *`num`* is set to a non-zero value, an IceGrid node will permit servers started by the node to run with super-user privileges. Note that you should not set this property unless the node uses a secure endpoint; otherwise, clients can start arbitrary processes with super-user privileges on the node's machine.

The default value of *`num`* is zero.

IceGrid.Node.CollocateRegistry

Synopsis

`IceGrid.Node.CollocateRegistry=`*`num`*

Copyright 2016, ZeroC, Inc.

Description

If *num* is set to a value larger than zero, the node collocates the IceGrid registry.

The collocated registry is configured with the same properties as the standalone IceGrid registry.

IceGrid.Node.Data

Synopsis

`IceGrid.Node.Data=`*path*

Description

Defines the path of the IceGrid node data directory. The node creates `distrib`, `servers`, and `tmp` subdirectories in this directory if they do not already exist. The `distrib` directory contains distribution files downloaded by the node from an IcePatch2 server. The `servers` directory contains configuration data for each deployed server. The `tmp` directory holds temporary files.

IceGrid.Node.DisableOnFailure

Synopsis

`IceGrid.Node.DisableOnFailure=`*num*

Description

The node considers a server to have terminated improperly if it has a non-zero exit code or if it exits due to one of the signals `SIGABRT`, `SIGBUS`, `SIGILL`, `SIGFPE`, or `SIGSEGV`. The node marks such a server as disabled if *num* is a non-zero value; a disabled server cannot be activated on demand. For values of *num* greater than zero, the server is disabled for *num* seconds. If *num* is a negative value, the server is disabled indefinitely, or until it is explicitly enabled or started via an administrative action. The default value is zero, meaning the node does not disable servers in this situation.

IceGrid.Node.Name

Synopsis

`IceGrid.Node.Name=`*name*

Description

Defines the *name* of the IceGrid node. All nodes using the same registry must have unique names; a node refuses to start if there is a node with the same name running already. This property must be defined for each node.

IceGrid.Node.Output

Synopsis

`IceGrid.Node.Output=`*path*

Description

Defines the path of the IceGrid node output directory. If set, the node redirects the `stdout` and `stderr` output of the started servers to files named `server.out` and `server.err` in this directory. Otherwise, the started servers share the `stdout` and `stderr` of the node's process.

IceGrid.Node.PrintServersReady

Synopsis

Copyright 2016, ZeroC, Inc.

`IceGrid.Node.PrintServersReady=`*`token`*

Description

The IceGrid node prints "*`token`* ready" on standard output after all the servers managed by the node are ready. This is useful for scripts that wish to wait until all servers have been started and are ready for use.

IceGrid.Node.ProcessorSocketCount

Synopsis

`IceGrid.Node.ProcessorSocketCount=`*`num`*

Description

This property sets the number of processor sockets. This value is reported by the `icegridadmin` `node sockets` command. On Windows Vista (or later), Windows Server 2008 (or later), and Linux systems, the number of processor sockets is set automatically by the Ice run time. On other systems, the run time cannot obtain the socket count from the operating system; you can use this property to set the number of processor sockets manually on such systems.

IceGrid.Node.PropertiesOverride

Synopsis

`IceGrid.Node.PropertiesOverride=`*`overrides`*

Description

Defines a list of properties that override the properties defined in server deployment descriptors. For example, in some cases it is desirable to set the property `Ice.Default.Host` for servers, but not in server deployment descriptors. The property definitions must be separated by white space.

IceGrid.Node.RedirectErrToOut

Synopsis

`IceGrid.Node.RedirectErrToOut=`*`num`*

Description

If *`num`* is set to a value larger than zero, the `stderr` of each started server is redirected to the server's `stdout`.

IceGrid.Node.Trace.Activator

Synopsis

`IceGrid.Node.Trace.Activator=`*`num`*

Description

The activator trace level:

0	No activator trace (default).
1	Trace process activation, termination.
2	Like 1, but more verbose, including process signaling and more diagnostic messages on process activation.
3	Like 2, but more verbose, including more diagnostic messages on process activation (e.g., path, working directory, and arguments of the activated process).

Copyright 2016, ZeroC, Inc.

IceGrid.Node.Trace.Adapter

Synopsis

`IceGrid.Node.Trace.Adapter=`*`num`*

Description

The object adapter trace level:

0	No object adapter trace (default).
1	Trace object adapter addition, removal.
2	Like 1, but more verbose, including object adapter activation and deactivation and more diagnostic messages.
3	Like 2, but more verbose, including object adapter transitional state change (for example, "waiting for activation").

IceGrid.Node.Trace.Patch

Synopsis

`IceGrid.Node.Trace.Patch=`*`num`*

Description

The patch trace level:

0	No patching trace (default).
1	Show summary of patch progress.
2	Like 1, but more verbose, including download statistics.
3	Like 2, but more verbose, including checksum information.

IceGrid.Node.Trace.Replica

Synopsis

`IceGrid.Node.Trace.Replica=`*`num`*

Description

The replica trace level:

0	No replica trace (default).
1	Trace session lifecycle between nodes and replicas.
2	Like 1, but more verbose, including session establishment attempts and failures.
3	Like 2, but more verbose, including keep alive messages sent to the replica.

IceGrid.Node.Trace.Server

Synopsis

`IceGrid.Node.Trace.Server=`*`num`*

Copyright 2016, ZeroC, Inc.

Description

The server trace level:

0	No server trace (default).
1	Trace server addition, removal.
2	Like 1, but more verbose, including server activation and deactivation, property updates, and more diagnostic messages.
3	Like 2, but more verbose, including server transitional state change (activating and deactivating).

IceGrid.Node.UserAccountMapper

Synopsis

`IceGrid.Node.UserAccountMapper=`*`proxy`*

Description

Specifies the proxy of an object that implements the `IceGrid::UserAccountMapper` interface for customizing the user accounts under which servers are activated. The IceGrid node invokes this proxy to map session identifiers (the user ID for sessions created with a user ID and password, or the distinguished name for sessions created from a secure connection) to user accounts.

As a proxy property, you can configure additional aspects of the proxy using properties.

IceGrid.Node.UserAccounts

Synopsis

`IceGrid.Node.UserAccounts=`*`file`*

Description

Specifies the file name of an IceGrid node user account map file. Each line of the file must contain an identifier and a user account, separated by white space. The identifier will be matched against the client session identifier (the user ID for sessions created with a user ID and password, or the distinguished name for sessions created from a secure connection). This user account map file is used by the node to map session identifiers to user accounts. This property is ignored if `IceGrid.Node.UserAccountMapper` is defined.

IceGrid.Node.WaitTime

Synopsis

`IceGrid.Node.WaitTime=`*`num`*

Description

Defines the interval in seconds that IceGrid waits for server activation and deactivation.

If a server is automatically activated and does not register its object adapter endpoints within this time interval, the node assumes there is a problem with the server and returns an empty set of endpoints to the client.

If a server is being gracefully deactivated and IceGrid does not detect the server deactivation during this time interval, IceGrid kills the server.

The default value is 60 seconds.

IceGrid.Registry.AdminCryptPasswords

Synopsis

Copyright 2016, ZeroC, Inc.

```
IceGrid.Registry.AdminCryptPasswords=file
```

Description

Specifies the file name of an IceGrid registry access control list for administrative clients. Each line of the file must contain a user name and a password, separated by white space. The password must a `MCF` encoded string as described here. If this property is not defined, the default value is `admin-passwords`. This property is ignored if `IceGrid.Registry.AdminPermissionsVerifier` is defined.

IceGrid.Registry.AdminPermissionsVerifier

Synopsis

```
IceGrid.Registry.AdminPermissionsVerifier=proxy
```

Description

Specifies the proxy of an object that implements the `Glacier2::PermissionsVerifier` interface for controlling access to IceGrid administrative sessions. The IceGrid registry invokes this proxy to validate each new administrative session created by a client with the `IceGrid::Registry` interface.

As a proxy property, you can configure additional aspects of the proxy using properties.

IceGrid.Registry.AdminSessionFilters

Synopsis

```
IceGrid.Registry.AdminSessionFilters=num
```

Description

This property controls whether IceGrid establishes filters for sessions created with the IceGrid session manager. If *num* is set to a value larger than zero, IceGrid establishes these filters, so Glacier2 limits access to the `IceGrid::AdminSession` object and the `IceGrid::Admin` object that is returned by the `getAdmin` operation. If *num* is set to zero, IceGrid does not establish filters, so access to these objects is controlled solely by Glacier2's configuration.

The default value is 1.

IceGrid.Registry.AdminSessionManager.*AdapterProperty*

Synopsis

```
IceGrid.Registry.AdminSessionManager.AdapterProperty=value
```

Description

The IceGrid registry uses the adapter name `IceGrid.Registry.AdminSessionManager` for the object adapter that processes incoming requests from IceGrid administrative sessions. Therefore, adapter properties can be used to configure this adapter. (Note any setting of `IceGrid.Registry.AdminSessionManager.AdapterId` is ignored because the registry always provides a direct adapter.)

For security reasons, defining endpoints for this object adapter is optional. If you do define endpoints, they should only be accessible to Glacier2 routers used to create IceGrid administrative sessions.

IceGrid.Registry.AdminSSLPermissionsVerifier

Synopsis

```
IceGrid.Registry.AdminSSLPermissionsVerifier=proxy
```

Description

Copyright 2016, ZeroC, Inc.

Specifies the proxy of an object that implements the `Glacier2::SSLPermissionsVerifier` interface for controlling access to IceGrid administrative sessions. The IceGrid registry invokes this proxy to validate each new administrative session created by a client from a secure connection with the `IceGrid::Registry` interface.

As a proxy property, you can configure additional aspects of the proxy using the properties.

IceGrid.Registry.Client.*AdapterProperty*

Synopsis

`IceGrid.Registry.Client.`*`AdapterProperty`*`=`*`value`*

Description

IceGrid uses the adapter name `IceGrid.Registry.Client` for the object adapter that processes incoming requests from clients. Therefore, adapter properties can be used to configure this adapter. (Note any setting of `IceGrid.Registry.Client.AdapterId` is ignored because the registry always provides a direct adapter.)

Note that `IceGrid.Registry.Client.Endpoints` controls the client endpoint for the registry. The port numbers 4061 (for TCP) and 4062 (for SSL) are reserved for the registry by the Internet Assigned Numbers Authority (IANA).

IceGrid.Registry.CryptPasswords

Synopsis

`IceGrid.Registry.CryptPasswords=`*`file`*

Description

Specifies the file name of an IceGrid registry access control list. Each line of the file must contain a user name and a password, separated by white space. The password must a `MCF` encoded string as described here. If this property is not defined, the default value is `passwords`. This property is ignored if `IceGrid.Registry.PermissionsVerifier` is defined.

IceGrid.Registry.Data

Synopsis

`IceGrid.Registry.Data=`*`path`*

Description

Defines the path of the IceGrid registry data directory. The directory specified in *`path`* must already exist.

IceGrid.Registry.DefaultTemplates

Synopsis

`IceGrid.Registry.DefaultTemplates=`*`path`*

Description

Defines the path name of an XML file containing default template descriptors. A sample file named `config/templates.xml` that contains convenient server templates for Ice services is provided in the Ice distribution.

IceGrid.Registry.Discovery.*AdapterProperty*

Synopsis

Copyright 2016, ZeroC, Inc.

`IceGrid.Registry.Discovery.`*`AdapterProperty`*`=`*`value`*

Description

The IceGrid registry creates an object adapter named `IceGrid.Registry.Discovery` for receiving multicast discovery queries from clients. If not otherwise defined by `IceGrid.Registry.Discovery.Endpoints`, the endpoint for this object adapter is composed as follows:

```
udp -h addr -p port [--interface intf]
```

where *`addr`* is the value of `IceGrid.Registry.Discovery.Address`, *`port`* is the value of `IceGrid.Registry.Discovery.Port`, and *`intf`* is the value of `IceGrid.Registry.Discovery.Interface`.

You don't normally need to set other properties for this object adapter.

IceGrid.Registry.Discovery.Address

Synopsis

`IceGrid.Registry.Discovery.Address=`*`addr`*

Description

Specifies the multicast IP address to use for receiving multicast discovery queries. If not defined, the default value depends on the setting of `Ice.IPv4`: if enabled (the default), IceDiscovery uses the address `239.255.0.1`, otherwise IceDiscovery assumes the application wants to use IPv6 and defaults to the address `ff15::1` instead. This property is used to compose the endpoint of the `IceGrid.Registry.Discovery` object adapter.

IceGrid.Registry.Discovery.Enabled

Synopsis

`IceDiscovery.Enabled=`*`num`*

Description

If *`num`* is a value larger than zero, the registry creates the `IceGrid.Registry.Discovery` object adapter and listens for multicast discovery queries. If not defined, the default value is `1`. Set this property to zero to disable multicast discovery.

IceGrid.Registry.Discovery.Interface

Synopsis

`IceGrid.Registry.Discovery.Interface=`*`intf`*

Description

Specifies the IP address of the interface to use for receiving multicast discovery queries. If not defined, the operating system will select a default interface to send and receive the UDP multicast datagrams. This property is used to compose the endpoint of the `IceGrid.Registry.Discovery` object adapter.

IceGrid.Registry.Discovery.Port

Synopsis

`IceGrid.Registry.Discovery.Port=`*`port`*

Description

Specifies the multicast port to use for receiving multicast discovery queries. If not set, the default value is `4061`. This property is used to

Copyright 2016, ZeroC, Inc.

compose the endpoint of the `IceGrid.Registry.Discovery` object adapter.

IceGrid.Registry.DynamicRegistration

Synopsis

`IceGrid.Registry.DynamicRegistration=`*`num`*

Description

If *`num`* is set to a value larger than zero, the locator registry does not require Ice servers to preregister object adapters and replica groups, but rather creates them automatically if they do not exist. If this property is not defined, or *`num`* is set to zero, an attempt to register an unknown object adapter or replica group causes adapter activation to fail with `Ice.NotRegisteredException`. An object adapter registers itself when the *`adapter`*`.AdapterId` property is defined. The *`adapter`*`.ReplicaGroupId` property identifies the replica group.

IceGrid.Registry.Internal.*AdapterProperty*

Synopsis

`IceGrid.Registry.Internal.`*`AdapterProperty`*`=`*`value`*

Description

The IceGrid registry uses the adapter name `IceGrid.Registry.Internal` for the object adapter that processes incoming requests from nodes and slave replicas. Therefore, adapter properties can be used to configure this adapter. (Note any setting of `IceGrid.Registry.Internal.AdapterId` is ignored because the registry always provides a direct adapter.)

IceGrid.Registry.NodeSessionTimeout

Synopsis

`IceGrid.Registry.NodeSessionTimeout=`*`num`*

Description

Each IceGrid node establishes a session with the registry that must be refreshed periodically. If a node does not refresh its session within *`num`* seconds, the node's session is destroyed and the servers deployed on that node become unavailable to new clients. If not specified, the default value is 30 seconds.

IceGrid.Registry.PermissionsVerifier

Synopsis

`IceGrid.Registry.PermissionsVerifier=`*`proxy`*

Description

Specifies the proxy of an object that implements the `Glacier2::PermissionsVerifier` interface for controlling access to IceGrid sessions. The IceGrid registry invokes this proxy to validate each new client session created by a client with the `IceGrid::Registry` interface.

As a proxy property, you can configure additional aspects of the proxy using properties.

IceGrid.Registry.ReplicaName

Synopsis

`IceGrid.Registry.ReplicaName=`*`name`*

Copyright 2016, ZeroC, Inc.

Description

Specifies the name of a registry replica. If not defined, the default value is `Master`, which is the name reserved for the master replica. Each registry replica must have a unique name.

IceGrid.Registry.ReplicaSessionTimeout

Synopsis

`IceGrid.Registry.ReplicaSessionTimeout=`*`num`*

Description

Each IceGrid registry replica establishes a session with the master registry that must be refreshed periodically. If a replica does not refresh its session within *`num`* seconds, the replica's session is destroyed and the replica no longer receives replication information from the master registry. If not specified, the default value is 30 seconds.

IceGrid.Registry.Server.*AdapterProperty*

Synopsis

`IceGrid.Registry.Server.`*`AdapterProperty=value`*

Description

The IceGrid registry uses the adapter name `IceGrid.Registry.Server` for the object adapter that processes incoming requests from servers. Therefore, adapter properties can be used to configure this adapter. (Note any setting of `IceGrid.Registry.Server.AdapterId` is ignored because the registry always provides a direct adapter.)

IceGrid.Registry.SessionFilters

Synopsis

`IceGrid.Registry.SessionFilters=`*`num`*

Description

This property controls whether IceGrid establishes filters for sessions created with the IceGrid session manager. If *`num`* is set to a value larger than zero, IceGrid establishes these filters, so Glacier2 limits access to the `IceGrid::Query` and `IceGrid::Session` objects, and to objects and adapters allocated by the session. If *`num`* is set to zero, IceGrid does not establish filters, so access to objects is controlled solely by Glacier2's configuration.

The default value is 0.

IceGrid.Registry.SessionManager.*AdapterProperty*

Synopsis

`IceGrid.Registry.SessionManager.`*`AdapterProperty=value`*

Description

The IceGrid registry uses the adapter name `IceGrid.Registry.SessionManager` for the object adapter that processes incoming requests from client sessions. Therefore, adapter properties can be used to configure this adapter. (Note any setting of `IceGrid.Registry.SessionManager.AdapterId` is ignored because the registry always provides a direct adapter.)

For security reasons, defining endpoints for this object adapter is optional. If you do define endpoints, they should only be accessible to Glacier2 routers used to create IceGrid client sessions.

Copyright 2016, ZeroC, Inc.

IceGrid.Registry.SessionTimeout

Synopsis

`IceGrid.Registry.SessionTimeout=`*`num`*

Description

IceGrid clients or administrative clients might establish a session with the registry. This session must be refreshed periodically. If the client does not refresh its session within *`num`* seconds, the session is destroyed. If not specified, the default value is 30 seconds.

IceGrid.Registry.SSLPermissionsVerifier

Synopsis

`IceGrid.Registry.SSLPermissionsVerifier=`*`proxy`*

Description

Specifies the proxy of an object that implements the `Glacier2::SSLPermissionsVerifier` interface for controlling access to IceGrid sessions. The IceGrid registry invokes this proxy to validate each new client session created by a client from a secure connection with the `IceGrid::Registry` interface.

As a proxy property, you can configure additional aspects of the proxy using properties.

IceGrid.Registry.Trace.Adapter

Synopsis

`IceGrid.Registry.Trace.Adapter=`*`num`*

Description

The object adapter trace level:

`0`	No object adapter trace (default).
`1`	Trace object adapter registration, removal, and replication.

IceGrid.Registry.Trace.Application

Synopsis

`IceGrid.Registry.Trace.Application=`*`num`*

Description

The application trace level:

0	No application trace (default).
1	Trace application addition, update, and removal.

IceGrid.Registry.Trace.Locator

Synopsis

`IceGrid.Registry.Trace.Locator=`*`num`*

Copyright 2016, ZeroC, Inc.

Description

The locator and locator registry trace level:

0	No locator trace (default).
1	Trace failures to locate an adapter or object, and failures to register adapter endpoints.
2	Like 1, but more verbose, including registration of adapter endpoints.

IceGrid.Registry.Trace.Node

Synopsis

`IceGrid.Registry.Trace.Node=`*`num`*

Description

The node trace level:

0	No node trace (default).
1	Trace node registration, removal.
2	Like 1, but more verbose, including load statistics.

IceGrid.Registry.Trace.Object

Synopsis

`IceGrid.Registry.Trace.Object=`*`num`*

Description

The object trace level:

0	No object trace (default).
1	Trace object registration, removal.

IceGrid.Registry.Trace.Patch

Synopsis

`IceGrid.Registry.Trace.Patch=`*`num`*

Description

The patch trace level:

0	No patching trace (default).
1	Show summary of patch progress.

IceGrid.Registry.Trace.Replica

Synopsis

Copyright 2016, ZeroC, Inc.

`IceGrid.Registry.Trace.Replica=`*`num`*

Description

The server trace level:

0	No server trace (default).
1	Trace session lifecycle between master replica and slaves.

IceGrid.Registry.Trace.Server

Synopsis

`IceGrid.Registry.Trace.Server=`*`num`*

Description

The server trace level:

0	No server trace (default).
1	Trace the addition and removal of servers in the Registry database.

IceGrid.Registry.Trace.Session

Synopsis

`IceGrid.Registry.Trace.Session=`*`num`*

Description

The session trace level:

0	No client or admin session trace (default).
1	Trace client or admin session registration, removal.
2	Like 1, but more verbose, includes keep alive messages.

IceGrid.Registry.UserAccounts

Synopsis

`IceGrid.Registry.UserAccounts=`*`file`*

Description

Specifies the file name of an IceGrid registry user account map file. Each line of the file must contain an identifier and a user account, separated by white space. The identifier will be matched against the client session identifier (the user ID for sessions created with a user ID and password, or the distinguished name for sessions created from a secure connection). This user account map file is used by IceGrid nodes to map session identifiers to user accounts if the nodes' `IceGrid.Node.UserAccountMapper` property is set to the proxy `IceGrid/RegistryUserAccountMapper`.

Copyright 2016, ZeroC, Inc.

IceGridAdmin.*

On this page:

- IceGridAdmin.AuthenticateUsingSSL
- IceGridAdmin.Discovery.Address
- IceGridAdmin.Discovery.Interface
- IceGridAdmin.Discovery.Lookup
- IceGridAdmin.Discovery.Reply.AdapterProperty
- IceGridAdmin.Host
- IceGridAdmin.InstanceName
- IceGridAdmin.Password
- IceGridAdmin.Port
- IceGridAdmin.Replica
- IceGridAdmin.Trace.Observers
- IceGridAdmin.Trace.SaveToRegistry
- IceGridAdmin.Username

IceGridAdmin.AuthenticateUsingSSL

Synopsis

`IceGridAdmin.AuthenticateUsingSSL=`*`num`*

Description

If *`num`* is a value greater than zero, `icegridadmin` uses SSL authentication when establishing its session with the IceGrid registry. If not defined or the value is zero, `icegridadmin` uses user name and password authentication.

IceGridAdmin.Discovery.Address

Synopsis

`IceGridAdmin.Discovery.Address=`*`addr`*

Description

Specifies the multicast IP address to use for sending multicast discovery queries. If not defined, the default value depends on the setting of `Ice.IPv4`: if enabled (the default), the client uses the address `239.255.0.1`, otherwise the client assumes it should use IPv6 and defaults to the address `ff15::1` instead. This property is used to compose the value of `IceGridAdmin.Discovery.Lookup`.

IceGridAdmin.Discovery.Interface

Synopsis

`IceGridAdmin.Discovery.Interface=`*`intf`*

Description

Specifies the IP address of the interface to use for sending multicast discovery queries. If not defined, the operating system will select a default interface to send the UDP multicast datagrams. This property is used to compose the value of `IceGridAdmin.Discovery.Lookup`.

IceGridAdmin.Discovery.Lookup

Synopsis

`IceGridAdmin.Discovery.Lookup=`*`endpoint`*

Description

Copyright 2016, ZeroC, Inc.

Specifies the endpoint that the client uses to send multicast discovery queries. If not defined, the endpoint is composed as follows:

```
udp -h addr -p port [--interface intf]
```

where *addr* is the value of `IceGridAdmin.Discovery.Address`, *port* is the value of `IceGridAdmin.Port`, and *intf* is the value of `IceGridAdmin.Discovery.Interface`.

IceGridAdmin.Discovery.Reply.*AdapterProperty*

Synopsis

`IceGridAdmin.Discovery.Reply.`*`AdapterProperty=value`*

Description

The client creates an object adapter named `IceGridAdmin.Discovery.Reply` for receiving replies to multicast discovery queries. If not otherwise defined by `IceGridAdmin.Discovery.Reply.Endpoints`, the endpoint for this object adapter is composed as follows:

```
udp [--interface intf]
```

where *intf* is the value of `IceGridAdmin.Discovery.Interface`. A fixed port is not necessary for this endpoint.

You don't normally need to set other properties for this object adapter.

IceGridAdmin.Host

Synopsis

`IceGridAdmin.Host=`*`host`*

Description

When used together with `IceGridAdmin.Port`, `icegridadmin` connects directly to the target registry at the specified host and port.

IceGridAdmin.InstanceName

Synopsis

`IceGridAdmin.InstanceName=`*`name`*

Description

Specifies the name of an IceGrid instance to which `icegridadmin` will connect.

When using multicast discovery, you can define this property to limit your discovery results only to those locators deployed for the given instance, in case you have multiple unrelated IceGrid instances deployed that use the same multicast address and port.

IceGridAdmin.Password

Synopsis

`IceGridAdmin.Password=`*`password`*

Description

Specifies the password that `icegridadmin` should use when authenticating its session with the IceGrid registry. For security reasons you may prefer not to define a password in a plain-text configuration property, in which case you should omit this property and allow `icegridadmin` to prompt you for it interactively. This property is ignored when SSL authentication is enabled via `IceGridAdmin.AuthenticateUsingSSL`.

Copyright 2016, ZeroC, Inc.

IceGridAdmin.Port

Synopsis

`IceGridAdmin.Port=`*`port`*

Description

When used together with `IceGridAdmin.Host`, `icegridadmin` connects directly to the target registry at the specified host and port.

When using multicast discovery, this property specifies the port to use for sending multicast discovery queries. This property is also used to compose the value of `IceGridAdmin.Discovery.Lookup`.

If not set, the default value is `4061`.

IceGridAdmin.Replica

Synopsis

`IceGridAdmin.Replica=`*`name`*

Description

Specifies the name of the registry replica that `icegridadmin` should contact. If not defined, the default value is `Master`.

IceGridAdmin.Trace.Observers

Synopsis

`IceGridAdmin.Trace.Observers=`*`num`*

Description

If *`num`* is a value greater than zero, the IceGrid graphical administrative client displays trace information about the observer callbacks it receives from the registry. If not defined, the default value is zero.

IceGridAdmin.Trace.SaveToRegistry

Synopsis

`IceGridAdmin.Trace.SaveToRegistry=`*`num`*

Description

If *`num`* is a value greater than zero, the IceGrid graphical administrative client displays trace information about the modifications it commits to the registry. If not defined, the default value is zero.

IceGridAdmin.Username

Synopsis

`IceGridAdmin.Username=`*`name`*

Description

Specifies the username that `icegridadmin` should use when authenticating its session with the IceGrid registry. This property is ignored when SSL authentication is enabled via `IceGridAdmin.AuthenticateUsingSSL`.

Copyright 2016, ZeroC, Inc.

IceLocatorDiscovery.*

This page describes the properties supported by the IceLocatorDiscovery plug-in.

On this page:

- IceLocatorDiscovery.Address
- IceLocatorDiscovery.InstanceName
- IceLocatorDiscovery.Interface
- IceLocatorDiscovery.Locator.AdapterProperty
- IceLocatorDiscovery.Lookup
- IceLocatorDiscovery.Port
- IceLocatorDiscovery.Reply.AdapterProperty
- IceLocatorDiscovery.RetryCount
- IceLocatorDiscovery.RetryDelay
- IceLocatorDiscovery.Timeout

IceLocatorDiscovery.Address

Synopsis

`IceLocatorDiscovery.Address=`*`addr`*

Description

Specifies the multicast IP address to use for sending multicast queries. If not defined, the default value depends on other property settings:

- If `Ice.PreferIPv6Address` is enabled or `Ice.IPv4` is disabled, IceLocatorDiscovery uses the IPv6 address `ff15::1`
- Otherwise IceLocatorDiscovery uses `239.255.0.1`

IceLocatorDiscovery.InstanceName

Synopsis

`IceLocatorDiscovery.InstanceName=`*`name`*

Description

Specifies the name of a locator instance. If you have multiple unrelated locators deployed that use the same multicast address and port, you can define this property to limit your discovery results only to those locators deployed for the given instance. If not defined, the plug-in adopts the instance name of the first locator to respond to a query; if a subsequent query discovers a locator with a different instance name, the plug-in logs a message and ignores the result.

IceLocatorDiscovery.Interface

Synopsis

`IceLocatorDiscovery.Interface=`*`intf`*

Description

Specifies the IP address of the interface to use for sending multicast queries. If not defined, the operating system will select a default interface to send the UDP multicast datagrams.

IceLocatorDiscovery.Locator.*AdapterProperty*

Synopsis

`IceLocatorDiscovery.Locator.`*`AdapterProperty=value`*

Copyright 2016, ZeroC, Inc.

Description

IceLocatorDiscovery creates an object adapter named `IceLocatorDiscovery.Locator`, therefore all of the object adapter properties can be set.

You don't normally need to set properties for this object adapter.

IceLocatorDiscovery.Lookup

Synopsis

`IceLocatorDiscovery.Lookup=`*`endpoint`*

Description

Specifies the multicast endpoint that a client uses to send discovery queries. If not defined, the endpoint is composed as follows:

```
udp -h addr -p port --interface intf
```

where *`addr`* is the value of `IceLocatorDiscovery.Address`, *`port`* is the value of `IceLocatorDiscovery.Port`, and *`intf`* is the value of `IceLocatorDiscovery.Interface`.

IceLocatorDiscovery.Port

Synopsis

`IceLocatorDiscovery.Port=`*`port`*

Description

Specifies the multicast port to use for sending multicast queries. If not set, the default value is `4061`.

IceLocatorDiscovery.Reply.*AdapterProperty*

Synopsis

`IceLocatorDiscovery.Reply.`*`AdapterProperty=value`*

Description

IceLocatorDiscovery creates an object adapter named `IceLocatorDiscovery.Reply` for receiving replies to multicast discovery queries. If not otherwise defined by `IceLocatorDiscovery.Reply.Endpoints`, the endpoint for this object adapter is composed as follows:

```
udp --interface intf
```

where *`intf`* is the value of `IceLocatorDiscovery.Interface`. A fixed port is not necessary for this endpoint.

You don't normally need to set other properties for this object adapter.

IceLocatorDiscovery.RetryCount

Synopsis

`IceLocatorDiscovery.RetryCount=`*`num`*

Description

Specifies the maximum number of times that the plug-in will retry sending UDP multicast queries before giving up. The `IceLocatorDiscovery.Timeout` property determines how long the plug-in waits for a reply before trying again. If not defined, the default retry count is 3.

Copyright 2016, ZeroC, Inc.

IceLocatorDiscovery.RetryDelay

Synopsis

`IceLocatorDiscovery.RetryDelay=`*`num`*

Description

If the plug-in fails to receive any responses to a query after retrying the number of times specified by `IceLocatorDiscovery.RetryCount`, the plug-in waits at least *`num`* milliseconds before starting another round of query attempts. If not defined, the default value is `2000`.

IceLocatorDiscovery.Timeout

Synopsis

`IceLocatorDiscovery.Timeout=`*`num`*

Description

Specifies the time interval in milliseconds to wait for replies to UDP multicast queries. If no server replies during this time interval, the client will retry the request the number of times specified by `IceLocatorDiscovery.RetryCount`. If not defined, the default timeout is `300`.

Copyright 2016, ZeroC, Inc.

IceMX.Metrics.*

On this page:

- IceMX.Metrics.view.Accept.attribute
- IceMX.Metrics.view.Disabled
- IceMX.Metrics.view.GroupBy
- IceMX.Metrics.view.Reject.attribute
- IceMX.Metrics.view.RetainDetached

Metrics view are configured with the properties described below. The *view* below can be replaced with one of the following:

- IceMX.Metrics.*view-name*
- IceMX.Metrics.*view-name*.Map.*map-name*
- IceMX.Metrics.*view-name*.Map.*map-name*.Map.*submap-name*

If a view is defined without Map properties, the view will contain all the metrics map known by the Metrics facet. If a view defines one or more map properties it will only contain these maps.

For a list of supported maps see:

- The Metrics Facet
- Glacier2 Metrics
- IceStorm Metrics

IceMX.Metrics.*view*.Accept.*attribute*

Synopsis

`IceMX.Metrics.`*`view`*`.Accept.`*`attribute`*`=`*`regexp`*

Description

This property defines a rule to accept the monitoring of an instrumented object or operation based on the value of one of its attribute. If the *attribute* matches the specified *regex* and if it satisfies other `Accept` and `Reject` filters the instrumented object or operation will be monitored.

For example, to accept monitoring instrumented objects or operations which are from the object adapter named "MyAdapter", you can setup the following accept property:

- `IceMX.Metrics.MyView.Accept.parent=MyAdapter`

IceMX.Metrics.*view*.Disabled

Synopsis

`IceMX.Metrics.`*`view`*`.Disabled=`*`num`*

Description

If *num* is set to a value larger than zero, the metrics view or the map is disabled. This property is useful to pre-configure a view or map. The view can be disabled initially to not incur overhead and enabled only when needed at runtime.

IceMX.Metrics.*view*.GroupBy

Synopsis

`IceMX.Metrics.`*`view`*`.GroupBy=`*`delimited attributes`*

Description

This property defines how metrics are grouped and how the ID of each metrics object is created. The grouping is based on attributes specific to the instrumented object or operation. For example, you can group the invocation metrics by operation name or proxy identity. All the

Copyright 2016, ZeroC, Inc.

invocations with the same operation name or proxy identity will record metrics using the same metrics object. You can specify several attributes to group metrics based on multiple attributes. You must delimit the attributes with delimiters when specify the value of the GroupBy property. A delimiter is any character which is not an alpha numeric or the dot character. Attributes which can be used to specify the value of this property are defined in relevant section of the Ice manual. Here are some examples of GroupBy properties.

- `IceMX.Metrics.MyView.GroupBy=operation`
- `IceMX.Metrics.MyView.GroupBy=identity [operation]`
- `IceMX.Metrics.MyView.GroupBy=remoteHost:remotePort`

IceMX.Metrics.*view*.Reject.*attribute*

Synopsis

`IceMX.Metrics.`*`view`*`.Reject.`*`attribute`*`=`*`regexp`*

Description

This property defines a rule to accept the monitoring of an instrumented object or operation based on the value of one of its attribute. If the *a`ttribute`* matches the specified *`regex`* and if it satisfies other `Accept` and `Reject` filters the instrumented object or operation will be monitored.

For example, to reject monitoring instrumented objects or operations which are from the object adapter named "Ice.Admin", you can setup the following reject property:

- `IceMX.Metrics.MyView.Reject.parent=Ice\.Admin`

IceMX.Metrics.*view*.RetainDetached

Synopsis

`IceMX.Metrics.`*`view`*`.RetainedDetached=`*`num`*

Description

If *`num`* is set to a value larger than zero, up to `num` metrics object whose `current` value is 0 will be kept in memory by the metrics map. This is useful to prevent indefinite memory growth if the monitoring of an instrumented object or operation creates a unique metrics object, only the last *`num`* metrics object will be kept in memory. The default value is `10`, meaning that at most ten metrics object with a `current` value equal to 0 will be retained by the metrics map.

Copyright 2016, ZeroC, Inc.

IcePatch2.*

On this page:

- IcePatch2.AdapterProperty
- IcePatch2.Directory
- IcePatch2.InstanceName

IcePatch2.*AdapterProperty*

Synopsis

`IcePatch2.`*`AdapterProperty`*`=`*`value`*

Description

IcePatch2 uses the adapter name `IcePatch2` for the server. Therefore, adapter properties can be used to configure this adapter.

IcePatch2.Directory

Synopsis

`IcePatch2.Directory=`*`dir`*

Description

The IcePatch2 server uses this property to determine the data directory if no data directory is specified on the command line.

This property is also used by IcePatch2 clients to determine the local data directory.

IcePatch2.InstanceName

Synopsis

`IcePatch2.InstanceName=`*`name`*

Description

Specifies the identity category for well-known IcePatch2 objects. If defined, the identity of the `IcePatch2::Admin` interface becomes *`name`*`/admin` and the identity of the `IcePatch2::FileServer` interface becomes *`name`*`/server`.

If not defined, the default value is `IcePatch2`.

Copyright 2016, ZeroC, Inc.

IcePatch2Client.*

On this page:

- IcePatch2Client.ChunkSize
- IcePatch2Client.Directory
- IcePatch2Client.Proxy
- IcePatch2Client.Remove
- IcePatch2Client.Thorough

IcePatch2Client.ChunkSize

Synopsis

`IcePatch2Client.ChunkSize=`*`kilobytes`*

Description

The IcePatch2 client uses this property to determine how many kilobytes are retrieved with each call to `getFileCompressed`.

The default value is `100`.

IcePatch2Client.Directory

Synopsis

`IcePatch2Client.Directory=`*`dir`*

Description

The IcePatch2 client uses this property to determine the local data directory.

IcePatch2Client.Proxy

Synopsis

`IcePatch2Client.Proxy=`*`proxy`*

Description

The IcePatch2 client uses this property to locate the IcePatch2 server.

IcePatch2Client.Remove

Synopsis

`IcePatch2Client.Remove=`*`num`*

Description

This property determines whether IcePatch2 clients delete files that exist locally, but not on the server. A negative or zero value prevents removal of files. A value of `1` enables removal and causes the client to halt with an error if removal of a file fails. A value of `2` or greater also enables removal, but causes the client to silently ignore errors during removal.

The default value is `1`.

IcePatch2Client.Thorough

Copyright 2016, ZeroC, Inc.

Synopsis

`IcePatch2Client.Thorough=`*`num`*

Description

This property determines whether IcePatch2 clients recompute checksums. Any value greater than zero is interpreted as true. The default value is `0` (false).

Copyright 2016, ZeroC, Inc.

IceSSL.*

On this page:

- IceSSL Property Overview
- IceSSL.Alias
- IceSSL.CAs
- IceSSL.CertFile
- IceSSL.CertStore
- IceSSL.CertStoreLocation
- IceSSL.CertVerifier
- IceSSL.CheckCertName
- IceSSL.CheckCRL
- IceSSL.Ciphers
- IceSSL.DefaultDir
- IceSSL.DH.bits
- IceSSL.DHParams
- IceSSL.EntropyDaemon
- IceSSL.FindCert
- IceSSL.InitOpenSSL
- IceSSL.Keychain
- IceSSL.KeychainPassword
- IceSSL.Keystore
- IceSSL.KeystorePassword
- IceSSL.KeystoreType
- IceSSL.Password
- IceSSL.PasswordCallback
- IceSSL.PasswordRetryMax
- IceSSL.Protocols
- IceSSL.ProtocolVersionMax
- IceSSL.ProtocolVersionMin
- IceSSL.Random
- IceSSL.Trace.Security
- IceSSL.TrustOnly
- IceSSL.TrustOnly.Client
- IceSSL.TrustOnly.Server
- IceSSL.TrustOnly.Server.AdapterName
- IceSSL.Truststore
- IceSSL.TruststorePassword
- IceSSL.TruststoreType
- IceSSL.UsePlatformCAs
- IceSSL.VerifyDepthMax
- IceSSL.VerifyPeer

IceSSL Property Overview

The IceSSL implementations for our supported platforms use many of the same configuration properties. However, there are some properties that are specific to certain platforms or languages. For properties with such limitations, we list the supported platforms or underlying SSL libraries in the synopsis and provide additional platform-specific notes if necessary. You'll see the following platforms, languages and SSL libraries listed in the property reference:

- SChannel (C++ on Windows)
- SecureTransport (C++ on OS X)
- OpenSSL (C++ on Linux)
- Java
- .NET
- iOS

A property is supported by all platforms if no limitations are mentioned.

Finally, note that Ice for Objective-C and the Ice extensions for Python, Ruby and PHP use IceSSL for C++, therefore they use the IceSSL properties for SChannel, SecureTransport or OpenSSL as appropriate for the target platform.

JavaScript and Universal Windows (UWP) Apps
These IceSSL properties have no effect on JavaScript and UWP apps.

Copyright 2016, ZeroC, Inc.

IceSSL.Alias

Synopsis

`IceSSL.Alias=`*`alias`* (Java)

Description

Selects a particular certificate from the key store specified by `IceSSL.Keystore`. The certificate identified by *`alias`* is presented to the peer request during authentication.

IceSSL.CAs

Synopsis

`IceSSL.CAs=`*`path`* (SChannel, SecureTransport, OpenSSL, .NET, iOS)

Description

Specifies the path name of a file containing the certificates of trusted certificate authorities (CAs).

If you wish to use the CA certificates bundled with your platform, leave this property unset and enable `IceSSL.UsePlatformCAs`.

Platform Notes

SChannel, SecureTransport, .NET

The file can be encoded using the DER or PEM formats. When using PEM, the file can contain multiple certificates.

IceSSL attempts to locate *`path`* as specified; if the given path is relative but does not exist, IceSSL also attempts to locate *`path`* relative to the default directory defined by `IceSSL.DefaultDir`.

OpenSSL

The file must be encoded using the PEM format and can contain multiple certificates. If *`path`* refers to a directory, the directory must be prepared in advance using the OpenSSL utility `c_rehash`.

IceSSL attempts to locate *`path`* as specified; if the given path is relative but does not exist, IceSSL also attempts to locate *`path`* relative to the default directory defined by `IceSSL.DefaultDir`.

iOS

The certificate must be encoded using the DER format.

If `IceSSL.DefaultDir` is also defined, IceSSL attempts to open the specified CA certificate file as `Resources/`*`DefaultDir/path`* in the application's resource bundle and as *`DefaultDir/path`* in the file system.

If `IceSSL.DefaultDir` is not defined, IceSSL attempts to open the specified CA certificate file as `Resources/`*`path`* in the application's resource bundle and as *`path`* in the file system.

If this property is not defined, IceSSL looks for suitable CA certificates in the user's keychains and in the system keychain.

Java

See `IceSSL.Truststore`.

IceSSL.CertFile

Synopsis

Copyright 2016, ZeroC, Inc.

`IceSSL.CertFile=`*`file`* (SecureTransport, .NET, iOS)
`IceSSL.CertFile=`*`file`*`[;`*`file`*`]` (SChannel)
`IceSSL.CertFile=`*`file`*`[:`*`file`*`]` (OpenSSL)

Description

Specifies a file that contains the program's certificate and the corresponding private key. The file name may be specified relative to the default directory defined by `IceSSL.DefaultDir`.

Platform Notes

SChannel

The file must use the PFX (PKCS#12) format and contain the certificate and its private key. If a password is required to load the file, the application must either install a password handler or supply the password using `IceSSL.Password`, otherwise IceSSL will reject the certificate.

This property accepts certificates for both RSA and DSA. To specify both certificates, separate the file names using the platform's path separator character.

IceSSL attempts to locate *`file`* as specified; if the given path is relative but does not exist, IceSSL also attempts to locate *`file`* relative to the default directory defined by `IceSSL.DefaultDir`.

SecureTransport

The file must use the PFX (PKCS#12) format and contain the certificate and its private key. If a password is required to load the file, OS X will use its default graphical password prompt unless the application has installed a password handler or supplied the password using `IceSSL.Password`. Define `IceSSL.Keychain` to import this certificate into the specified keychain.

IceSSL attempts to locate *`file`* as specified; if the given path is relative but does not exist, IceSSL also attempts to locate *`file`* relative to the default directory defined by `IceSSL.DefaultDir`.

OpenSSL

The file must use the PFX (PKCS#12) format and contain the certificate and its private key. If a password is required to load the file, OpenSSL will prompt the user at the terminal unless the application has installed a password handler or supplied the password using `IceSSL.Password`.

This property accepts certificates for both RSA and DSA. To specify both certificates, separate the file names using the platform's path separator character.

IceSSL attempts to locate *`file`* as specified; if the given path is relative but does not exist, IceSSL also attempts to locate *`file`* relative to the default directory defined by `IceSSL.DefaultDir`.

.NET

The file must use the PFX (PKCS#12) format and contain the certificate and its private key. The password for the file must be supplied using `IceSSL.Password`.

IceSSL attempts to locate *`file`* as specified; if the given path is relative but does not exist, IceSSL also attempts to locate *`file`* relative to the default directory defined by `IceSSL.DefaultDir`.

iOS

The file must use the PFX (PKCS#12) format and contain the certificate and its private key. The password for the file must be supplied using `IceSSL.Password`. The certificate is imported into the keychain identified by the `IceSSL.Keychain` property.

If `IceSSL.DefaultDir` is also defined, IceSSL attempts to open the specified certificate file as `Resources/`*`DefaultDir/file`* in the application's resource bundle and as *`DefaultDir/file`* in the file system.

If `IceSSL.DefaultDir` is not defined, IceSSL attempts to open the specified certificate file as `Resources/`*`file`* in the application's resource bundle and as *`file`* in the file system.

Java

See `IceSSL.Keystore`.

Copyright 2016, ZeroC, Inc.

IceSSL.CertStore

Synopsis

`IceSSL.CertStore=`*name* (SChannel, .NET)

Description

Specifies the name of a certificate store to use when locating certificates via `IceSSL.FindCert`. Legal values for *name* include `AddressBook`, `AuthRoot`, `CertificateAuthority`, `Disallowed`, `My`, `Root`, `TrustedPeople`, and `TrustedPublisher`. You can also use an arbitrary value for *name*.

If not specified, the default value is `My`.

IceSSL.CertStoreLocation

Synopsis

`IceSSL.CertStoreLocation=CurrentUser|LocalMachine` (SChannel, .NET)

Description

Specifies the location of a certificate store to use when locating certificates via `IceSSL.FindCert`. If not specified, the default value is `CurrentUser`.

An Ice program running as a Windows service will typically need to set this property to `LocalMachine`.

IceSSL.CertVerifier

Synopsis

`IceSSL.CertVerifier=`*classname* (Java, .NET)

Description

Specifies the name of a Java or .NET class that implements the `IceSSL.CertificateVerifier` interface for performing application-defined certificate verification.

Platform Notes

SChannel, SecureTransport, OpenSSL

C++ applications can install a certificate verifier programmatically.

IceSSL.CheckCertName

Synopsis

`IceSSL.CheckCertName=`*num*

Description

If *num* is a value greater than zero, IceSSL attempts to match the server's host name as specified in the proxy endpoint against the common name component of the server certificate's subject name. If no match is found, IceSSL attempts to match the host name against the DNS and IP address fields of the server certificate's subject alternative name. The search does not issue any DNS queries but simply performs a case-insensitive string match. The server's certificate is accepted if its common name or any of its DNS or IP addresses matches the host

Copyright 2016, ZeroC, Inc.

name in the proxy endpoint. IceSSL skips this validation step if the server does not supply a certificate, or if the proxy endpoint does not include a host name and `Ice.Default.Host` is not defined. This property has no affect on a server's validation of a client's certificate. If no match is found, IceSSL aborts the connection attempt and raises an exception. If not defined, the default value is zero.

IceSSL.CheckCRL

Synopsis

`IceSSL.CheckCRL=`*`num`* (.NET)

Description

If *num* is a value greater than zero, IceSSL checks the certificate revocation list (CRL) to determine if the peer's certificate has been revoked. The value for *num* determines the resulting behavior:

`0`	Disables CRL checking.
`1`	If a certificate is revoked, IceSSL aborts the connection, logs a message and raises an exception. If a certificate's revocation status is unknown, IceSSL logs a message but accepts the certificate.
`2`	If a certificate is revoked or its revocation status is unknown, IceSSL aborts the connection, logs a message and raises an exception.

The `IceSSL.Trace.Security` property must be set to a non-zero value to see CRL-related log messages. If `IceSSL.CheckCRL` is not defined, the default value is zero.

IceSSL.Ciphers

Synopsis

`IceSSL.Ciphers=`*`ciphers`* (SChannel, SecureTransport, OpenSSL, Java)

Description

Specifies the cipher suites that IceSSL is allowed to negotiate. A cipher suite is a set of algorithms that satisfies the four requirements for establishing a secure connection: signing and authentication, key exchange, secure hashing, and encryption. Some algorithms satisfy more than one requirement, and there are many possible combinations.

If not specified, the plug-in uses the security provider's default cipher suites. Enable `IceSSL.Trace.Security` and carefully review the application's log output to determine which cipher suites are enabled by default, or to verify your cipher suite configuration.

Platform Notes

SChannel

The value of this property is a whitespace-separated list that can include any of the following:

- `3DES`
- `AES_128`
- `AES_256`
- `DES`
- `RC2`
- `RC4`

For example, the following setting enables AES cipher suites:

```
IceSSL.Ciphers=AES_128 AES_256
```

SecureTransport

The property value is interpreted as a list of tokens delimited by white space. The plug-in executes the tokens in the order of appearance in order to assemble the list of enabled cipher suites. The table below describes the tokens:

Copyright 2016, ZeroC, Inc.

`ALL`	Reserved keyword that enables all supported cipher suites. If specified, it must be the first token in the list. Use with caution as it may enable low-security cipher suites.
`NONE`	Reserved keyword that disables all cipher suites. If specified, it must be the first token in the list. Use `NONE` to start with an empty set of cipher suites and then add only those suites you want to allow.
NAME	Includes the cipher suite whose name matches *NAME* exactly.
`!`*NAME*	Excludes the cipher suite whose name matches *NAME* exactly.
`(`*EXP*`)`	Includes all cipher suites whose names contain the given regular expression *EXP*.
`!(`*EXP*`)`	Excludes all cipher suites whose names contain the given regular expression *EXP*.

For example, the following setting disables all cipher suites except those that support 256-bit AES with SHA256:

```
IceSSL.Ciphers=NONE (AES_256.*SHA256)
```

Note that no warning is given if an unrecognized cipher is specified.

OpenSSL

The value of this property is passed directly to the OpenSSL library and the list of supported ciphers depends on how your installation of OpenSSL was compiled. You can obtain a complete list of the supported cipher suites using the command `openssl ciphers`. This command will likely generate a long list. To simplify the selection process, OpenSSL supports several classes of ciphers. Classes and ciphers can be excluded by prefixing them with an exclamation point. The special keyword `@STRENGTH` sorts the cipher list in order of their strength, so that SSL gives preference to the more secure ciphers when negotiating a cipher suite. The `@STRENGTH` keyword must be the last element in the list. The classes are:

`ALL`	Enables all supported cipher suites. This class should be used with caution, as it may enable low-security cipher suites.
`ADH`	Anonymous ciphers.
`LOW`	Low bit-strength ciphers.
`EXP`	Export-crippled ciphers.

Here is an example of a reasonable setting:

```
IceSSL.Ciphers=ALL:!ADH:!LOW:!EXP:!MD5:@STRENGTH
```

This value excludes the ciphers with low bit-strength and known problems, and orders the remaining ciphers according to their strength. Note that no warning is given if an unrecognized cipher is specified.

Java

The property value is interpreted as a list of tokens delimited by white space. The plug-in executes the tokens in the order of appearance in order to assemble the list of enabled cipher suites. The table below describes the tokens:

`ALL`	Reserved keyword that enables all supported cipher suites. If specified, it must be the first token in the list. Use with caution as it may enable low-security cipher suites.
`NONE`	Reserved keyword that disables all cipher suites. If specified, it must be the first token in the list. Use `NONE` to start with an empty set of cipher suites and then add only those suites you want to allow.
NAME	Enables the cipher suite matching the given name.

Copyright 2016, ZeroC, Inc.

!*NAME*	Disables the cipher suite matching the given name.
(EXP)	Enables cipher suites whose names contain the regular expression *EXP*. For example, the value `NONE (.*DH_anon.*AES.*)` selects only cipher suites that use anonymous Diffie-Hellman authentication with AES encryption.
!(*EXP*)	Disables cipher suites whose names contain the regular expression *EXP*. For example, the value `ALL !(.*DH_anon.*AES.*)` enables all cipher suites except those that use anonymous Diffie-Hellman authentication with AES encryption.

IceSSL.DefaultDir

Synopsis

`IceSSL.DefaultDir=`*`path`*

Description

Specifies the default directory in which to look for certificates, key stores, and other files. See the descriptions of the relevant properties for more information.

IceSSL.DH.*bits*

Synopsis

`IceSSL.DH.`*`bits`*`=`*`file`* (OpenSSL)

Description

Specifies a *`file`* containing Diffie Hellman parameters whose key length is *`bits`*, as shown in the following example:

```
IceSSL.DH.1024=dhparams1024.pem
```

IceSSL supplies default parameters for key lengths of 512, 1024, 2048, and 4096 bits, which are used if no user-defined parameters of the desired key length are specified. The parameters must be encoded using the PEM format.

IceSSL attempts to locate *`file`* as specified; if the given path is relative but does not exist, IceSSL also attempts to locate *`file`* relative to the default directory defined by `IceSSL.DefaultDir`.

Platform Notes

SChannel

Anonymous Diffie Hellman ciphers are not supported on Windows.

SecureTransport

See `IceSSL.DHParams`.

IceSSL.DHParams

Synopsis

`IceSSL.DHParams=`*`file`* (SecureTransport)

Description

Copyright 2016, ZeroC, Inc.

Specifies a *file* containing Diffie Hellman parameters. The parameters must be encoded using the DER format.

This property only affects server (incoming) connections. Clients obtain their DH parameters when negotiating a SSL/TLS connection with the server.

IceSSL attempts to locate *file* as specified; if the given path is relative but does not exist, IceSSL also attempts to locate *file* relative to the default directory defined by `IceSSL.DefaultDir`.

If this property is not specified, OS X will generate its own set of Diffie Hellman parameters for the process. Computing these parameters at run time can take up to 30 seconds, so we recommend generating them in advance and defining this property.

You can generate Diffie Hellman parameters using the `openssl dhparam` command.

Platform Notes

SChannel

Anonymous Diffie Hellman ciphers are not supported on Windows.

OpenSSL

Use `IceSSL.DH.bits`.

IceSSL.EntropyDaemon

Synopsis

`IceSSL.EntropyDaemon=file` (OpenSSL)

Description

Specifies a Unix domain socket for the entropy gathering daemon, from which OpenSSL gathers entropy data to initialize its random number generator.

IceSSL.FindCert

Synopsis

`IceSSL.FindCert=criteria` (SChannel, SecureTransport, .NET)

Description

Builds a collection of certificates that will be used for authentication.

A server requires a certificate for authentication purposes, therefore IceSSL selects the first certificate in the accumulated collection. This is normally the certificate loaded via `IceSSL.CertFile`, if that property was defined. Otherwise, IceSSL selects one of the certificates identified by `IceSSL.FindCert`.

Platform Notes

SChannel, .NET

IceSSL queries a certificate store for matching certificates and adds them to the application's certificate collection. The settings for `IceSSL.CertStore` and `IceSSL.CertStoreLocation` determine the target certificate store to be queried.

The value for *criteria* may be `*`, in which case all of the certificates in the store are selected. Otherwise, *criteria* must be one or more *field*`:`*value* pairs separated by white space. The valid field names are described below:

Copyright 2016, ZeroC, Inc.

`Issuer`	Matches a substring of the issuer's name.
`IssuerDN`	Matches the issuer's entire distinguished name.
`Serial`	Matches the certificate's serial number.
`Subject`	Matches a substring of the subject's name.
`SubjectDN`	Matches the subject's entire distinguished name.
`SubjectKeyId`	Matches the certificate's subject key identifier.
`Thumbprint`	Matches the certificate's thumbprint.

The field names are case-insensitive. If multiple criteria are specified, only certificates that match all criteria are selected. Values must be enclosed in single or double quotes to preserve white space.

SecureTransport

IceSSL queries the keychain for matching certificates and adds them to the application's certificate collection. IceSSL uses the keychain identified in `IceSSL.Keychain`, or the user's default keychain if `IceSSL.Keychain` is not defined.

The value for *`criteria`* must be one or more *`field`*`:`*`value`* pairs separated by white space. The valid field names are described below:

`Label`	Matches the user-visible label.
`Serial`	Matches the certificate's serial number.
`Subject`	Matches a substring of the subject's name.
`SubjectKeyId`	Matches the certificate's subject key identifier.

The field names are case-insensitive. If multiple criteria are specified, only certificates that match all criteria are selected. Values must be enclosed in single or double quotes to preserve white space.

Java

Use `IceSSL.Alias`.

IceSSL.InitOpenSSL

Synopsis

`IceSSL.InitOpenSSL=`*`num`* (OpenSSL)

Description

Indicates whether IceSSL should perform the global initialization tasks for the OpenSSL library. The default value is 1, meaning IceSSL will initialize OpenSSL. An application can set this property to zero if it wishes to perform the OpenSSL initialization itself, which can be useful when the application uses multiple components that depend on OpenSSL.

IceSSL.Keychain

Synopsis

`IceSSL.Keychain=`*`name`* (SecureTransport, iOS)

Description

Specifies the name of a keychain in which to import the certificate identified by `IceSSL.CertFile`. Set `IceSSL.KeychainPassword` if the specified keychain has a password.

Platform Notes

Copyright 2016, ZeroC, Inc.

SecureTransport

A relative path name is opened relative to the current working directory. If the specified keychain file does not exist, a new file is created. If not defined, IceSSL uses the user's default keychain.

iOS

If not defined, the keychain named `login` is used by default. Note that this property is only relevant for the iPhone simulator and OS X targets.

IceSSL.KeychainPassword

Synopsis

`IceSSL.KeychainPassword=`*`password`* (SecureTransport, iOS)

Description

Specifies the password for the keychain identified by `IceSSL.Keychain`. If not defined, IceSSL attempts to open the keychain without a password.

Platform Notes

SecureTransport

If not defined and the keychain requires a password, OS X will ask the user in a graphical prompt.

iOS

This property is only relevant for the iPhone simulator and OS X targets.

IceSSL.Keystore

Synopsis

`IceSSL.Keystore=`*`file`* (Java)

Description

Specifies a key store file containing certificates and their private keys. If the key store contains multiple certificates, you should specify a particular one to use for authentication using `IceSSL.Alias`. IceSSL first attempts to open *`file`* as a class loader resource and then as a regular file. If the given path is relative but does not exist, IceSSL also attempts to locate it relative to the default directory defined by `IceSSL.DefaultDir`. The format of the file is determined by `IceSSL.KeystoreType`.

If this property is not defined, the application will not be able to supply a certificate during SSL handshaking. As a result, the application may not be able to negotiate a secure connection, or might be required to use an anonymous cipher suite.

IceSSL.KeystorePassword

Synopsis

`IceSSL.KeystorePassword=`*`password`* (Java)

Description

Specifies the password used to verify the integrity of the key store defined by `IceSSL.Keystore`. The integrity check is skipped if this property is not defined.

It is a security risk to use a plain-text password in a configuration file.

Copyright 2016, ZeroC, Inc.

IceSSL.KeystoreType

Synopsis

`IceSSL.KeystoreType=`*`type`* (Java)

Description

Specifies the format of the key store file defined by `IceSSL.Keystore`. Legal values are `JKS` and `PKCS12`. If not defined, the JVM's default value is used (normally `JKS`).

IceSSL.Password

Synopsis

`IceSSL.Password=`*`password`*

Description

Specifies the password necessary to decrypt the private key.

It is a security risk to use a plain-text password in a configuration file.

Platform Notes

SChannel, SecureTransport, OpenSSL

This property supplies the password that was used to secure the private key contained in the file defined by `IceSSL.CertFile`.

Java

This property supplies the password that was used to secure the private key contained in the key store defined by `IceSSL.Keystore`. All of the keys in the key store must use the same password.

.NET

This property supplies the password that was used to secure the file defined by `IceSSL.CertFile`.

iOS

This property supplies the password that was used to secure the file defined by `IceSSL.CertFile`.

IceSSL.PasswordCallback

Synopsis

`IceSSL.PasswordCallback=`*`classname`* (Java, .NET)

Description

Specifies the name of a Java or .NET class that implements the `IceSSL.PasswordCallback` interface. Using a password callback is a more secure alternative to specifying a password in a plain-text configuration file.

Platform Notes

SChannel, SecureTransport, OpenSSL

Use `setPasswordPrompt` to install a password callback in the plug-in.

Copyright 2016, ZeroC, Inc.

IceSSL.PasswordRetryMax

Synopsis

`IceSSL.PasswordRetryMax=`*`num`* (SChannel, SecureTransport, OpenSSL)

Description

Specifies the number of attempts IceSSL should allow the user to make when entering a password. If not defined, the default value is `3`.

IceSSL.Protocols

Synopsis

`IceSSL.Protocols=`*`list`* (SChannel, SecureTransport, OpenSSL, Java, .NET)

Description

Specifies the protocols to allow during SSL handshaking. Legal values are `SSL3`, `TLS1`, `TLS1_0` (alias for `TLS1`), `TLS1_1`, and `TLS1_2`. You may also specify multiple values, separated by commas or white space. If this property is not defined, the default setting is as follows:

```
IceSSL.Protocols=TLS1_0, TLS1_1, TLS1_2
```

In C#, the options `TLS1_1` and `TLS1_2` require .NET 4.5 or later.

As of Ice 3.6, SSLv3 is disabled by default.

Platform Notes

SecureTransport

Use `IceSSL.ProtocolVersionMin` and `IceSSL.ProtocolVersionMax`.

IceSSL.ProtocolVersionMax

Synopsis

`IceSSL.ProtocolVersionMax=`*`prot`* (SecureTransport)

Description

Specifies the maximum protocol to allow during SSL handshaking. Legal values are `SSL3`, `TLS1`, `TLS1_0` (alias for `TLS1`), `TLS1_1`, and `TLS1_2`. If this property is not defined, the default value is `TLS1_0`.

Platform Notes

SChannel, OpenSSL, Java, .NET

Use `IceSSL.Protocols`.

IceSSL.ProtocolVersionMin

Synopsis

Copyright 2016, ZeroC, Inc.

`IceSSL.ProtocolVersionMin=`*`prot`* (SecureTransport)

Description

Specifies the minimum protocol to allow during SSL handshaking. Legal values are `SSL3`, `TLS1`, `TLS1_0` (alias for `TLS1`), `TLS1_1`, and `TLS1_2`. If this property is not defined, the platform's default is used.

Platform Notes

SChannel, OpenSSL, Java, .NET

Use `IceSSL.Protocols`.

IceSSL.Random

Synopsis

`IceSSL.Random=`*`filelist`* (OpenSSL, Java)

Description

Specifies one or more files containing data to use when seeding the random number generator. The file names should be separated using the platform's path separator.

Platform Notes

OpenSSL

IceSSL attempts to locate each file as specified; if a given path is relative but does not exist, IceSSL also attempts to locate it relative to the default directory defined by `IceSSL.DefaultDir`.

Java

IceSSL first attempts to open each file as a class loader resource and then as a regular file. If a given path is relative but does not exist, IceSSL also attempts to locate it relative to the default directory defined by `IceSSL.DefaultDir`.

IceSSL.Trace.Security

Synopsis

`IceSSL.Trace.Security=`*`num`* (SChannel, SecureTransport, OpenSSL, Java, .NET)

Description

The SSL plug-in trace level:

`0`	No security tracing (default).
`1`	Display diagnostic information about SSL connections.

IceSSL.TrustOnly

Synopsis

`IceSSL.TrustOnly=`*`ENTRY`*`[;`*`ENTRY`*`;...]` (SChannel, SecureTransport, OpenSSL, Java, .NET)

Description

Identifies trusted and untrusted peers. This family of properties provides an additional level of authentication by using the peer certificate's distinguished name (DN) to decide whether to accept or reject a connection.

Copyright 2016, ZeroC, Inc.

Each *ENTRY* in the property value consists of relative distinguished name (RDN) components, formatted according to the rules in RFC 2253. Specifically, the components must be separated by commas, and any component that contains a comma must be escaped or enclosed in quotes. For example, the following two property definitions are equivalent:

```
IceSSL.TrustOnly=O="Acme, Inc.",OU=Sales
IceSSL.TrustOnly=O=Acme\, Inc.,OU="Sales"
```

Use a semicolon to separate multiple entries in a property:

```
IceSSL.TrustOnly=O=Acme\, Inc.,OU=Sales;O=Acme\, Inc.,OU=Marketing
```

By default, each entry represents an acceptance entry. A `!` character appearing at the beginning of an entry signifies a rejection entry. The order of the entries in a property is not important.

After the SSL engine has successfully completed its authentication process, IceSSL evaluates the relevant `IceSSL.TrustOnly` properties in an attempt to find an entry that matches the peer certificate's DN. For a match to be successful, the peer DN must contain an exact match for all of the RDN components in an entry. An entry may contain as many RDN components as you wish, depending on how narrowly you need to restrict access. The order of the RDN components in an entry is not important.

The connection semantics are described below:

1. IceSSL aborts the connection if any rejection or acceptance entries are defined and the peer does not supply a certificate.
2. IceSSL aborts the connection if the peer DN matches any rejection entry. (This is true even if the peer DN also matches an acceptance entry.)
3. IceSSL accepts the connection if the peer DN matches any acceptance entry, or if no acceptance entries are defined.

Our original example limits access to people in the sales and marketing departments:

```
IceSSL.TrustOnly=O=Acme\, Inc.,OU=Sales;O=Acme\, Inc.,OU=Marketing
```

If it later becomes necessary to deny access to certain individuals in these departments, you can add a rejection entry and restart the program:

```
IceSSL.TrustOnly=O=Acme\, Inc.,OU=Sales; O=Acme\, Inc.,OU=Marketing; !O=Acme\,
Inc.,CN=John Smith
```

While testing your trust configuration, you may find it helpful to set the `IceSSL.Trace.Security` property to a non-zero value, which causes IceSSL to display the DN of each peer during connection establishment.

This property affects incoming and outgoing connections. IceSSL also supports similar properties that affect only incoming connections or only outgoing connections.

IceSSL.TrustOnly.Client

Synopsis

`IceSSL.TrustOnly.Client=`*`ENTRY`*`[;`*`ENTRY`*`;...]` (SChannel, SecureTransport, OpenSSL, Java, .NET)
`IceSSL.TrustOnly.Client=`*`ID`* (iOS)

Description

Identifies trusted and untrusted peers for outgoing (client) connections. The entries defined in this property are combined with those of `IceSSL.TrustOnly`.

Platform Notes

Copyright 2016, ZeroC, Inc.

iOS

For an outgoing connection to succeed, the peer certificate's subject key identifier must match the property value exactly. The property value is formatted as a series of hexadecimal values separated by colons or spaces, as shown in the example below:

```
C2:E8:D3:33:D7:83:99:6E:08:F7:C2:34:31:F7:1E:8E:44:87:38:57
```

Since this value is specific to a single certificate authority, this property is valid only when used in conjunction with `IceSSL.CertAuthFile`.

IceSSL.TrustOnly.Server

Synopsis

`IceSSL.TrustOnly.Server=`*`ENTRY`*`[;`*`ENTRY`*`;...]` (SChannel, SecureTransport, OpenSSL, Java, .NET)

Description

Identifies trusted and untrusted peers for incoming ("server") connections. The entries defined in this property are combined with those of `IceSSL.TrustOnly`. To configure trusted and untrusted peers for a particular object adapter, use `IceSSL.TrustOnly.Server.AdapterName`.

IceSSL.TrustOnly.Server.*AdapterName*

Synopsis

`IceSSL.TrustOnly.Server.`*`AdapterName`*`=`*`ENTRY`*`[;`*`ENTRY`*`;...]` (SChannel, SecureTransport, OpenSSL, Java, .NET)

Description

Identifies trusted and untrusted peers for incoming (server) connections to the object adapter *`AdapterName`*. The entries defined in this property are combined with those of `IceSSL.TrustOnly` and `IceSSL.TrustOnly.Server`.

IceSSL.Truststore

Synopsis

`IceSSL.Truststore=`*`file`* (Java)

Description

Specifies a key store file containing the certificates of trusted certificate authorities. IceSSL first attempts to open *`file`* as a class loader resource and then as a regular file. If the given path is relative but does not exist, IceSSL also attempts to locate it relative to the default directory defined by `IceSSL.DefaultDir`. The format of the file is determined by `IceSSL.TruststoreType`.

If this property is not defined, IceSSL uses the value of `IceSSL.Keystore` by default. If no truststore is specified and the keystore does not contain a valid certificate chain, the application will not be able to authenticate the peer's certificate during SSL handshaking. As a result, the application may not be able to negotiate a secure connection, or might be required to use an anonymous cipher suite.

Platform Notes

SChannel, SecureTransport, OpenSSL, .NET

Use `IceSSL.CAs`.

IceSSL.TruststorePassword

Copyright 2016, ZeroC, Inc.

Synopsis

`IceSSL.TruststorePassword=`*`password`* (Java)

Description

Specifies the password used to verify the integrity of the key store defined by `IceSSL.Truststore`. The integrity check is skipped if this property is not defined.

It is a security risk to use a plain-text password in a configuration file.

Platform Notes

SChannel, SecureTransport, OpenSSL, .NET

Use `IceSSL.Password`.

IceSSL.TruststoreType

Synopsis

`IceSSL.TruststoreType=`*`type`* (Java)

Description

Specifies the format of the key store file defined by `IceSSL.Truststore`. Legal values are `JKS` and `PKCS12`. If not defined, the default value is `JKS`.

IceSSL.UsePlatformCAs

Synopsis

`IceSSL.UsePlatformCAs=`*`num`*

Description

If *`num`* is a value greater than zero, IceSSL uses the platform's bundled Root Certificate Authorities. This setting is ignored if `IceSSL.CAs` is defined.

If not defined, the default value is zero.

IceSSL.VerifyDepthMax

Synopsis

`IceSSL.VerifyDepthMax=`*`num`* (SChannel, SecureTransport, OpenSSL, Java, .NET)

Description

Specifies the maximum depth of a trusted peer's certificate chain, including the peer's certificate. A value of zero accepts chains of any length. If not defined, the default value is 3.

IceSSL.VerifyPeer

Synopsis

`IceSSL.VerifyPeer=`*`num`* (SChannel, SecureTransport, OpenSSL, Java, .NET)

Copyright 2016, ZeroC, Inc.

Description

Specifies the verification requirements to use during SSL handshaking. The legal values are shown in the table below. If this property is not defined, the default value is `2`.

`0`	For an outgoing connection, the client verifies the server's certificate (if an anonymous cipher is not used) but does not abort the connection if verification fails. For an incoming connection, the server does not request a certificate from the client.
`1`	For an outgoing connection, the client verifies the server's certificate and aborts the connection if verification fails. For an incoming connection, the server requests a certificate from the client and verifies it if one is provided, aborting the connection if verification fails.
`2`	For an outgoing connection, the semantics are the same as for the value `1`. For an incoming connection, the server requires a certificate from the client and aborts the connection if verification fails.

Platform Notes

.NET

This property has no effect on outgoing connections, since .NET always uses the semantics of value `2`. For an incoming connection, the value `0` has the same semantics as the value `1`.

Copyright 2016, ZeroC, Inc.

IceStorm Properties

IceStorm is an IceBox service that you can install using any name you like. For example:

```
IceBox.Service.DataFeed=IceStormService,...
```

The service name you choose is also used as the prefix for IceStorm's configuration properties. In the example above, the IceStorm configuration properties would use the `DataFeed` prefix, as in `DataFeed.Discard.Interval=10`.

In the property descriptions below, replace *`service`* with the service name from your IceStorm configuration.

On this page:

- service.Discard.Interval
- service.Election.ElectionTimeout
- service.Election.MasterTimeout
- service.Election.ResponseTimeout
- service.Flush.Timeout
- service.InstanceName
- service.Node.AdapterProperty
- service.NodeId
- service.Nodes.id
- service.Publish.AdapterProperty
- service.ReplicatedPublishEndpoints
- service.ReplicatedTopicManagerEndpoints
- service.Send.Timeout
- service.Send.QueueSizeMax
- service.Send.QueueSizeMaxPolicy
- service.TopicManager.AdapterProperty
- service.Trace.Election
- service.Trace.Replication
- service.Trace.Subscriber
- service.Trace.Topic
- service.Trace.TopicManager
- service.Transient

service.Discard.Interval

Synopsis

`service``.Discard.Interval=`*`num`*

Description

An IceStorm server detects when a subscriber to which it forwards events becomes non-functional and, at that point, stops delivery attempts to that subscriber for *`num`* seconds before trying to forward events to that subscriber again. The default value of this property is 60 seconds.

service.Election.ElectionTimeout

Synopsis

`service``.Election.ElectionTimeout=`*`num`*

Description

This property is used by a replicated IceStorm deployment. It specifies the interval in seconds at which a coordinator attempts to form larger groups of replicas. If not defined, the default value is 10.

Copyright 2016, ZeroC, Inc.

service.Election.MasterTimeout

Synopsis

`service.Election.MasterTimeout=num`

Description

This property is used by a replicated IceStorm deployment. It specifies the interval in seconds at which a slave checks the status of the coordinator. If not defined, the default value is 10.

service.Election.ResponseTimeout

Synopsis

`service.Election.ResponseTimeout=num`

Description

This property is used by a replicated IceStorm deployment. It specifies the interval in seconds that a replica waits for replies to an invitation to form a larger group. Lower priority replicas wait for intervals inversely proportional to the maximum priority:

```
ResponseTimeout + ResponseTimeout * (max - pri)
```

If not defined, the default value is 10.

service.Flush.Timeout

Synopsis

`service.Flush.Timeout=num`

Description

Defines the interval in milliseconds with which events are sent to batch subscribers. The default is 1000ms.

service.InstanceName

Synopsis

`service.InstanceName=name`

Description

Specifies an alternate identity category for all objects hosted by the IceStorm object adapters. If not specified, the default identity category is `IceStorm`.

service.Node.*AdapterProperty*

Copyright 2016, ZeroC, Inc.

Synopsis

service.Node.*AdapterProperty*=*value*

Description

In a replicated deployment, IceStorm uses the adapter name *service*.Node for the replica node's object adapter. Therefore, adapter properties can be used to configure this adapter.

service.NodeId

Synopsis

service.NodeId=*value*

Description

Specifies the node ID of an IceStorm replica, where *value* is a non-negative integer. The node ID is also used as the replica's priority, such that a larger value assigns higher priority to the replica. The replica with the highest priority becomes the coordinator of its group. This property must be defined for each replica.

service.Nodes.*id*

Synopsis

service.Nodes.*id*=*value*

Description

This property is used for a manual deployment of highly available IceStorm, in which each of the replicas must be explicitly configured with the proxies of all other replicas. The value is a proxy for the replica with the given node *id*. A replica's object identity has the form *instance-name*/node*id*, such as DemoIceStorm/node2.

service.Publish.*AdapterProperty*

Synopsis

service.Publish.*AdapterProperty*=*value*

Description

IceStorm uses the adapter name *service*.Publish for the object adapter that processes incoming requests from publishers. Therefore, a dapter properties can be used to configure this adapter.

service.ReplicatedPublishEndpoints

Synopsis

service.ReplicatedPublishEndpoints=*value*

Description

This property is used for a manual deployment of highly available IceStorm. It specifies the set of endpoints returned for the publisher proxy

Copyright 2016, ZeroC, Inc.

returned from `IceStorm::Topic::getPublisher`.

If this property is not defined, the publisher proxy returned by a topic instance points directly at that replica and, should the replica become unavailable, publishers will not transparently failover to other replicas.

service.ReplicatedTopicManagerEndpoints

Synopsis

`service.ReplicatedTopicManagerEndpoints=value`

Description

This property is used for a manual deployment of highly available IceStorm. It specifies the set of endpoints used in proxies that refer to a replicated topic. This set of endpoints should contain the endpoints of each IceStorm replica.

For example, the operation `IceStorm::TopicManager::create` returns a proxy that contains this set of endpoints.

service.Send.Timeout

Synopsis

`service.Send.Timeout=num`

Description

IceStorm applies a send timeout when it forwards events to subscribers. The value of this property determines how long IceStorm will wait for forwarding of an event to complete. If an event cannot be forwarded within *num* milliseconds, the subscriber is considered dead and its subscription is cancelled. The default value is 60 seconds. Setting this property to a negative value disables timeouts.

service.Send.QueueSizeMax

Synopsis

`service.Send.QueueSizeMax=num  (since Ice 3.6.3)`

Description

The value of this property determines how many events can be queued for a subscriber by IceStorm. When the maximum size is reached, the old events will either be dropped or the subscriber will be removed. Setting this property to a negative value specifies an infinite queue size. The default value is -1.

service.Send.QueueSizeMaxPolicy

Synopsis

`service.Send.QueueSizePolicy=RemoveSubscriber|DropEvents  (since Ice 3.6.3)`

Description

The value of this property specifies how IceStorm will behave if the maximum queue size is reached for a subscriber. If set to `RemoveSubscriber`, IceStorm will remove the subscriber as soon as the limit is reached. If set to `DropEvents`, older events will be removed to make room for new events. The default value is `RemoveSubscriber`.

Copyright 2016, ZeroC, Inc.

service.TopicManager.*AdapterProperty*

Synopsis

`service.TopicManager.AdapterProperty=value`

Description

IceStorm uses the adapter name `service.TopicManager` for the topic manager's object adapter. Therefore, adapter properties can be used to configure this adapter.

service.Trace.Election

Synopsis

`service.Trace.Election=num`

Description

Trace activity related to elections:

0	No election trace (default).
1	Trace election activity.

service.Trace.Replication

Synopsis

`service.Trace.Replication=num`

Description

Trace activity related to replication:

0	No replication trace (default).
1	Trace replication activity.

service.Trace.Subscriber

Synopsis

`service.Trace.Subscriber=num`

Description

The subscriber trace level:

0	No subscriber trace (default).

Copyright 2016, ZeroC, Inc.

1	Trace topic diagnostic information on subscription and unsubscription.
2	Like 1, but more verbose, including state transitions for a subscriber (such as going offline after a temporary network failure, and going online again after a successful retry, etc.).

service.Trace.Topic

Synopsis

`service.Trace.Topic=num`

Description

The topic trace level:

0	No topic trace (default).
1	Trace topic links, subscription, and unsubscription.
2	Like 1, but more verbose, including QoS information, and other diagnostic information.

service.Trace.TopicManager

Synopsis

`service.Trace.TopicManager=num`

Description

The topic manager trace level:

0	No topic manager trace (default).
1	Trace topic creation.

service.Transient

Synopsis

`service.Transient=num`

Description

If *num* is a value greater than zero, IceStorm runs in a fully transient mode in which no database is required. Replication is not supported in this mode. If not defined, the default value is zero.

Copyright 2016, ZeroC, Inc.

IceStormAdmin.*

On this page:

- IceStormAdmin.Host
- IceStormAdmin.Port
- IceStormAdmin.TopicManager.Default
- IceStormAdmin.TopicManager.name

IceStormAdmin.Host

Synopsis

`IceStormAdmin.Host=`*`host`*

Description

When used together with `IceStormAdmin.Port`, `icestormadmin` uses the `Finder` interface to discover the topic manager at the specified host and port.

> `icestormadmin` ignores this setting if you define one or more `IceStormAdmin.TopicManager` properties.

IceStormAdmin.Port

Synopsis

`IceStormAdmin.Port=`*`port`*

Description

When used together with `IceStormAdmin.Host`, `icestormadmin` uses the `Finder` interface to discover the topic manager at the specified host and port.

> `icestormadmin` ignores this setting if you define one or more `IceStormAdmin.TopicManager` properties.

IceStormAdmin.TopicManager.Default

Synopsis

`IceStormAdmin.TopicManager.Default=`*`proxy`*

Description

Defines the proxy for the default IceStorm topic manager. This property is used by `icestormadmin`. IceStorm applications may choose to use this property for their configuration as well.

IceStormAdmin.TopicManager.*name*

Synopsis

`IceStormAdmin.TopicManager.`*`name`*`=`*`proxy`*

Description

Defines a proxy for an IceStorm topic manager for `icestormadmin`. Properties with this pattern are used by `icestormadmin` if multiple

Copyright 2016, ZeroC, Inc.

topic managers are in use, for example:

```
IceStormAdmin.TopicManager.A=A/TopicManager:tcp -h x -p 9995
IceStormAdmin.TopicManager.B=Foo/TopicManager:tcp -h x -p 9995
IceStormAdmin.TopicManager.C=Bar/TopicManager:tcp -h x -p 9987
```

This sets the proxies for three topic managers. Note that *name* need not match the instance name of the corresponding topic manager — *name* simply serves as a tag. With these property settings, the `icestormadmin` commands that accept a topic can now specify a topic manager other than the default topic manager that is configured with `IceStormAdmin.TopicManager.Default`. For example:

```
current Foo
create myTopic
create Bar/myOtherTopic
```

This sets the current topic manager to the one with instance name `Foo`; the first `create` command then creates the topic within that topic manager, whereas the second `create` command uses the topic manager with instance name `Bar`.

Copyright 2016, ZeroC, Inc.

Ice Release Notes

The release notes provide information about an Ice release, including descriptions of significant new features and changes, instructions for upgrading from an earlier release, and important platform-specific details.

Topics

- Supported Platforms for Ice 3.6.3
- New Features in Ice 3.6
- Backward Compatibility of Ice Versions
- Upgrading your Application from Ice 3.6.0 or later
- Upgrading your Application from Ice 3.5
- Upgrading your Application from Ice 3.4
- Upgrading your Application from Ice 3.3
- Upgrading your Application from Ice 3.2 or Earlier Releases
- Platform Notes for Ice 3.6
- Using the Windows Binary Distribution
- Using the Linux Binary Distributions
- Using the macOS Binary Distribution
- Using Ice on Android
- Using the JavaScript Distribution
- Using the Python Distribution
- Using the Ruby Distribution
- Getting Started with Ice on AWS
- Building Third Party Packages from Source

Copyright 2016, ZeroC, Inc.

Supported Platforms for Ice 3.6.3

Ice 3.6.3 is supported on the platform, compiler, and environment combinations shown below. Other platforms and compilers might work as well but have not been tested.

On this page:

- Ice for C++
- Ice for Java
- Ice for .NET
- Ice for Universal Windows Platform (UWP)
- Ice for WinRT
- Ice for PHP
- Ice for Python
- Ice for Ruby
- Ice for JavaScript

Ice for C++

Platform	Compiler	Architecture
Windows 7 Windows 8.1 Windows 10 Windows Server 2008 R2 Windows Server 2012	Visual Studio 2012, Visual Studio 2013, Visual Studio 2015	x86, x64
Red Hat Enterprise Linux 6 Red Hat Enterprise Linux 7 Amazon Linux 2016.09 SuSE Linux Enterprise Server 11 SuSE Linux Enterprise Server 12 Ubuntu 14.04 (Trusty Tahr) Ubuntu 16.04 (Xenial Xerus)	GCC 4.4 GCC 4.8 GCC 4.8 GCC 4.3 GCC 4.8 GCC 4.8 GCC 5.3	i686, x86_64 i686, x86_64 x86_64 i586, x86_64 x86_64 i386, amd64 amd64
OS X 10.11 El Capitan macOS 10.12 Sierra	Xcode 7.3, Xcode 8 Xcode 8	x86_64 x86_64

Ice for Java

Platform	Environment
All Ice for C++ platforms	JDK 1.7, JDK 1.8
Android 4.2 or later	JDK 1.7, JDK 1.8, Android Studio

Ice for .NET

Platform	Compiler	Architecture
Windows 7 Windows 8.1 Windows 10 Windows Server 2008 R2 Windows Server 2012	Visual Studio 2012, Visual Studio 2013, Visual Studio 2015	x86, x64

Ice for Universal Windows Platform (UWP)

Platform	Compiler	Architecture

Copyright 2016, ZeroC, Inc.

Windows 10	Visual Studio 2015 (C++ only)	x86, x64

Ice for WinRT

Platform	Compiler	Architecture
Windows 8.1	Visual Studio 2013 (C++ only)	x86, x64

Ice for PHP

Platform	Environment	Architecture
Windows 7 Windows 8.1 Windows 10 Windows Server 2008 R2 Windows Server 2012	PHP 5.6	x86
Red Hat Enterprise Linux 6 Red Hat Enterprise Linux 7 Amazon Linux 2016.09 SuSE Linux Enterprise Server 11 SuSE Linux Enterprise Server 12 Ubuntu 14.04 (Trusty Tahr) Ubuntu 16.04 (Wily Werewolf)	PHP 5.3 PHP 5.4 PHP 5.3 PHP 5.3 PHP 5.5 PHP 5.5 PHP 7.0	i686, x86_64 i686, x86_64 x86_64 i586, x86_64 x86_64 i386, amd64 amd64
OS X 10.11 El Capitan macOS 10.12 Sierra	PHP 5.5 PHP 5.6	x86_64 x86_64

Ice for Python

Platform	Environment	Architecture
Windows 7 Windows 8.1 Windows 10 Windows Server 2008 R2 Windows Server 2012	Python 2.7, Python 3.4, Python 3.5	x86, x64
Red Hat Enterprise Linux 6 Red Hat Enterprise Linux 7 Amazon Linux 2016.09 SuSE Linux Enterprise Server 11 SuSE Linux Enterprise Server 12 Ubuntu 14.04 (Trusty Tahr) Ubuntu 16.04 (Xenial Xerus)	Python 2.6, Python 2.7, Python 3.4, Python 3.5	i686, x86_64 i686, x86_64 x86_64 i586, x86_64 x86_64 i386, amd64 amd64
OS X 10.11 El Capitan macOS 10.12 Sierra	Python 2.7, Python 3.4, Python 3.5	x86_64

Ice for Ruby

Platform	Environment	Architecture
Windows 7 Windows 8.1 Windows 10 Windows Server 2008 R2 Windows Server 2012	Ruby 2.3	x86, x64

Copyright 2016, ZeroC, Inc.

Red Hat Enterprise Linux 6 Red Hat Enterprise Linux 7 Amazon Linux 2016.09 SuSE Linux Enterprise Server 11 SuSE Linux Enterprise Server 12 Ubuntu 14.04 (Trusty Tahr) Ubuntu 16.04 (Xenial Xerus)	Ruby 1.8 to Ruby 2.3	i686, x86_64 i686, x86_64 x86_64 i586, x86_64 x86_64 i386, amd64 amd64
OS X 10.11 El Capitan macOS 10.12 Sierra	Ruby 2.0 to Ruby 2.3	x86_64

Ice for JavaScript

Platform	Environment
Web browser	Recent versions of Internet Explorer, Chrome, Firefox, Safari
Windows 7 Windows 8.1 Windows 10 Windows Server 2008 R2 Windows Server 2012 Ubuntu 14.04 (Trusty Tahr) Ubuntu 16.04 (Xenial Xerus) OS X 10.11 El Capitan macOS 10.12 Sierra	Node.js 0.10, Node.js 0.12, Node.js 4.0, Node.js 5.x Node.js 6.x

Copyright 2016, ZeroC, Inc.

New Features in Ice 3.6

This page describes notable additions and improvements in Ice 3.6. For a detailed list of the changes in this release, please refer to the changelog in the source tree. Our upgrade guide documents the changes that may affect the operation of your applications or have an impact on your source code.

On this page:

- New Features
  - JavaScript Mapping and WebSocket Transports
  - Objective-C Mapping
  - Multicast Discovery
    - IceDiscovery
    - IceLocatorDiscovery
  - IceGrid Custom Load Balancing
  - Invocation Timeouts
  - Invocation Cancellation
  - Active Connection Management
  - SSL
- Other New Features
  - Remote Logging
  - Garbage Collection Improvements
  - Transport Improvements
    - HTTP Proxy
    - Source Address
  - Batch Invocations
  - Collocated Invocations
  - Finder Interfaces
  - Java Interrupts
  - Java Lambda Functions
  - Java Buffers
  - C++ View Types
  - Linking with C++ Libraries on Windows
  - C++ Plug-in for Crypt Password Files
  - Support for -fvisibility=hidden with GCC and clang

New Features

Ice 3.6 includes a number of significant new features, many of them in response to requests from our commercial users.

JavaScript Mapping and WebSocket Transports

Ice for JavaScript offers a native and compact implementation of the Ice run time that you can incorporate into your web apps and gain seamless access to back-end Ice services. Our native JavaScript implementation of the Ice run time features a compact footprint, complete support for all Slice data types, compatibility with modern browsers as well as NodeJS, and an intuitive API to get you productive as quickly as possible.

Here's a really simple but complete example to give you an idea of what it's like to make a remote invocation using Ice in JavaScript:

Copyright 2016, ZeroC, Inc.

JavaScript

```
var communicator = Ice.initialize();
var proxy = communicator.stringToProxy("hello:wss -h www.myhost.com -p 443");
var hello = HelloPrx.uncheckedCast(proxy);
hello.sayHello().then(
    function() { console.log("sayHello done!"); }
).exception(
    function(ex) { console.log("something went wrong!"); }
).finally(
    function() { return communicator.destroy(); }
);
```

Experienced JavaScript developers might recognize our API's use of *promises*, which is a common solution for managing asynchronous activities. Our promise implementation follows existing conventions and offers a tremendous amount of power and convenience for simplifying your application code.

There have been a few changes to the JavaScript mapping since the preview release, as described in the upgrade guide.

Please refer to the manual for complete details on the JavaScript mapping.

JavaScript clients running in the browser use the industry-standard WebSocket protocol. They can talk directly to Ice servers listening on WebSocket endpoints. In our example above, the endpoint `wss -h www.myhost.com -p 443` directs the client to use a secure WebSocket connection via the standard HTTPS port. The WebSocket protocol is supported by all Ice language mappings.

The use of WebSockets isn't useful just for JavaScript clients running in the browser -- it can also be used by regular Ice clients as well. The WebSocket protocol can indeed greatly ease network deployment by allowing the use of the standard HTTP or HTTPS port as well as your existing Web infrastructure. For instance, the client can connect to:

- an Ice server listening directly on port 80 (HTTP) or port 443 (HTTPS),
- an HTTP server (e.g., Nginx) that supports reverse proxying of WebSocket connections,
- an HTTP load balancer (e.g., HAProxy) that supports forwarding WebSocket connections.

Here's how it works when using an HTTP server reverse proxy:

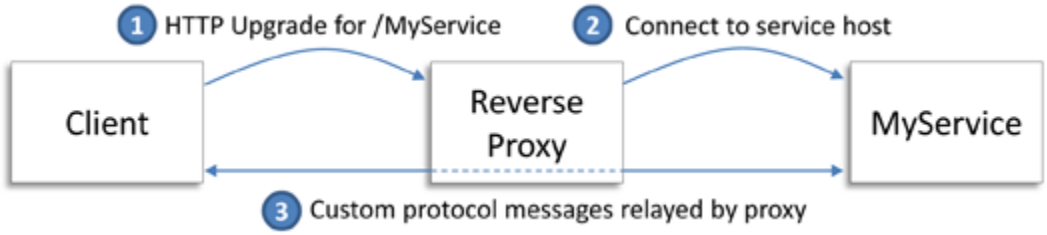

The client connects to the HTTP server and sends an HTTP request to upgrade the existing connection to support WebSocket. The message includes a resource path, such as `/MyService` in the example above, that the HTTP server translates into a target back-end service. The HTTP server establishes its own connection to the service and then relays messages between the client and the service. The service is an Ice server listening on a WS or WSS endpoint.

The WebSocket protocol provides an alternative to opening a port on your corporate firewall. You can instead leverage your existing web infrastructure to allow Ice clients on the other side of your firewall to access your Ice services.

Objective-C Mapping

The Objective-C mapping from Ice Touch is now included in the Ice distribution.

It currently lacks support for the following features:

Copyright 2016, ZeroC, Inc.

- Asynchronous method dispatch (AMD)

Please refer to the manual for complete details on the Objective-C mapping.

Multicast Discovery

In an effort to simplify bootstrapping and minimize configuration, we're introducing two new components that facilitate the discovery of Ice servers via UDP multicast.

IceDiscovery

IceDiscovery provides a lightweight location service that enables clients to discover objects at run time. Implemented as an Ice plug-in, the service requires only minimal configuration and very few changes to your application code. IceDiscovery is an ideal solution for applications that don't need the additional features that IceGrid offers, such as server activation and centralized administration, but could still take advantage of the benefits provided by indirect proxies and replication.

Consider the following code example:

Java

```
Ice.ObjectPrx proxy = communicator.stringToProxy("hello");
HelloPrx hello = HelloPrxHelper.checkedCast(proxy);
hello.sayHello();
```

Notice here that the code translates the stringified proxy `"hello"` containing no endpoint information. Ice resolves such proxies using a locator supplied by the IceDiscovery plug-in, all we need to do is configure the client to load the plug-in:

```
# Client configuration
Ice.Plugin.IceDiscovery=IceDiscovery:IceDiscovery.PluginFactory
```

At the proxy's first remote invocation (the call to `checkedCast` in the example above), IceDiscovery issues a multicast query to determine if any servers are available that provide the `hello` object. If more than one server responds, it means the object is replicated and the client selects one of the servers based on additional configuration settings. This discovery process is handled internally by the plug-in and is transparent to clients and servers.

UDP multicast is only used by IceDiscovery to locate available servers. Clients and servers communicate using the transport of your choosing.

In practical terms, using the plug-in means a client no longer needs to hard-code the addressing information of its servers in its source code or configuration files. Assuming the underlying network supports multicast, IceDiscovery enables servers to start, stop, and migrate to different hosts without affecting the ability of clients to find them, and without requiring any administrative changes.

IceLocatorDiscovery

In previous versions of Ice, using IceGrid was the primary way for applications to leverage the benefits of indirect proxies; the only other alternative was to write a custom "locator" service. Indirect proxies help to minimize coupling between clients and servers by eliminating most of the addressing information from application code and configuration files. The one exception has always been the `Ice.Default.Locator` property, which needed to contain at least one endpoint for an IceGrid registry:

```
# Client configuration before IceLocatorDiscovery
Ice.Default.Locator=IceGrid/Locator:tcp -h registryhost -p 4061
```

In a replicated IceGrid deployment consisting of a master replica and one or more slave replicas, the `Ice.Default.Locator` property would normally include endpoints for some or all of the replicas.

Copyright 2016, ZeroC, Inc.

The IceLocatorDiscovery plug-in alleviates the need to statically configure locator endpoints by using UDP multicast to dynamically discover the active locators in a network. Integrating the plug-in simply requires updating client configuration files to install the plug-in and remove the `Ice.Default.Locator` property. No changes are necessary in your application code.

IceGrid Custom Load Balancing

We've enhanced the IceGrid registry to support pluggable load balancing strategies, making it possible for you to implement application-specific load balancing algorithms or filter the registry's results based on user input.

Two kinds of filters are supported:

- Replica group filter
  The registry invokes a replica group filter each time a client requests the endpoints of a replica group or object adapter, as well as for calls to `findAllReplicas`. The registry passes information about the query that the filter can use in its implementation, including the list of object adapters participating in the replica group whose nodes are active at the time of the request. The object adapter list is initially ordered using the load balancing type configured for the replica group; the filter can modify this list however it chooses.
- Type filter
  The registry invokes a type filter for each query that a client issues to find a well-known object by type using the operations `findObjectByType`, `findAllObjectsByType`, and `findObjectByTypeOnLeastLoadedNode`. Included in the information passed to the filter is a list of proxies for the matching well-known objects; the filter implementation decides which of these proxies are returned to the client.

Filters are implemented in C++ as regular Ice plug-ins that you install in the IceGrid registry via configuration properties.

Invocation Timeouts

Timeout semantics have changed for the better with the addition of invocation timeouts. Ice still supports the existing timeout facility, referred to as a connection timeout or endpoint timeout, but now it only affects network activities. The new invocation timeout facility enables you to set an arbitrary timeout per invocation, independent of the connection timeout. Furthermore, when an invocation timeout occurs, it only affects that invocation and the connection remains open. The code below demonstrates the new semantics:

C++

```
Ice::ObjectPrx proxy = communicator->stringToProxy(...);
proxy = proxy->ice_timeout(2000);           // Connection timeout: 2sec
proxy = proxy->ice_invocationTimeout(5000); // Invocation timeout: 5sec
try {
    proxy->ice_ping();
} catch (const Ice::InvocationTimeoutException& ex) {
    // server's busy?
} catch (const Ice::TimeoutException& ex) {
    // network problem!
}
```

Review the upgrade guide to learn how your existing application might be impacted by and benefit from these changes.

Invocation Cancellation

It's now possible to cancel an asynchronous invocation using the `cancel` method of the `Ice::AsyncResult` object. This method prevents a queued invocation from being sent or, if the invocation has already been sent, ignores a reply if the server sends one. It has no effect on the server. A canceled invocation is considered to be completed and the result of the invocation is an `Ice::InvocationCanceledException`.

Active Connection Management

Copyright 2016, ZeroC, Inc.

We've enhanced the Active Connection Management facility to give applications more control over ACM behavior. Included in these enhancements is a new "heartbeat" feature that you can use in your applications when you need to ensure that a connection remains open. For example, you can use heartbeats to keep a Glacier2 session alive or prevent a bidirectional connection from being closed prematurely. This feature eliminates the need for an application to implement its own background thread to keep a connection or session alive. We've also added the ability to register a callback with the connection that is notified when the connection closes or receives a heartbeat. The code below shows how to enable heartbeats and install a callback on a connection:

C++

```
class CallbackI : public Ice::ConnectionCallback
{
    virtual void heartbeat(const Ice::ConnectionPtr& connection)
    {
        std::cout << "received heartbeat for connection:\n" <<
connection->toString() << std::endl;
    }

    virtual void closed(const Ice::ConnectionPtr& connection)
    {
         std::cout << "connection closed:\n" << connection->toString() <<
std::endl;
    }
};
Ice::ObjectPrx proxy = communicator->stringToProxy(...);
Ice::ConnectionPtr connection proxy->ice_getConnection();
connection->setACM(5, Ice::CloseOnInvocationAndIdle, Ice::HeartbeatAlways);
connection->setCallback(new CallbackI());
```

Refer to our upgrade guide for more information on how the ACM changes can impact existing Ice applications.

SSL

We've essentially re-written the IceSSL plug-in to make use of a platform's native SSL APIs where possible. On OS X, the plug-in now uses Apple's Secure Transport facility, and on Windows it now uses SChannel. Linux is the only platform on which IceSSL still uses OpenSSL.

By using native SSL APIs instead of OpenSSL for secure communications on OS X and Windows, your applications can more easily take advantage of any hotfixes provided by the operating system vendors when security issues arise.

If you use IceSSL, we recommend reviewing the upgrade guide because there have been some changes to the IceSSL APIs and configuration properties.

Other New Features

This section describes several more additions and improvements in Ice 3.6 that we'd like to bring to your attention.

Remote Logging

The Ice administrative facility adds a new logger facet that allows remote access to a program's log activity. The IceGrid administrative tools now use this feature to let you view the recent log output of Ice servers, IceBox services, and IceGrid nodes and registries.

Copyright 2016, ZeroC, Inc.

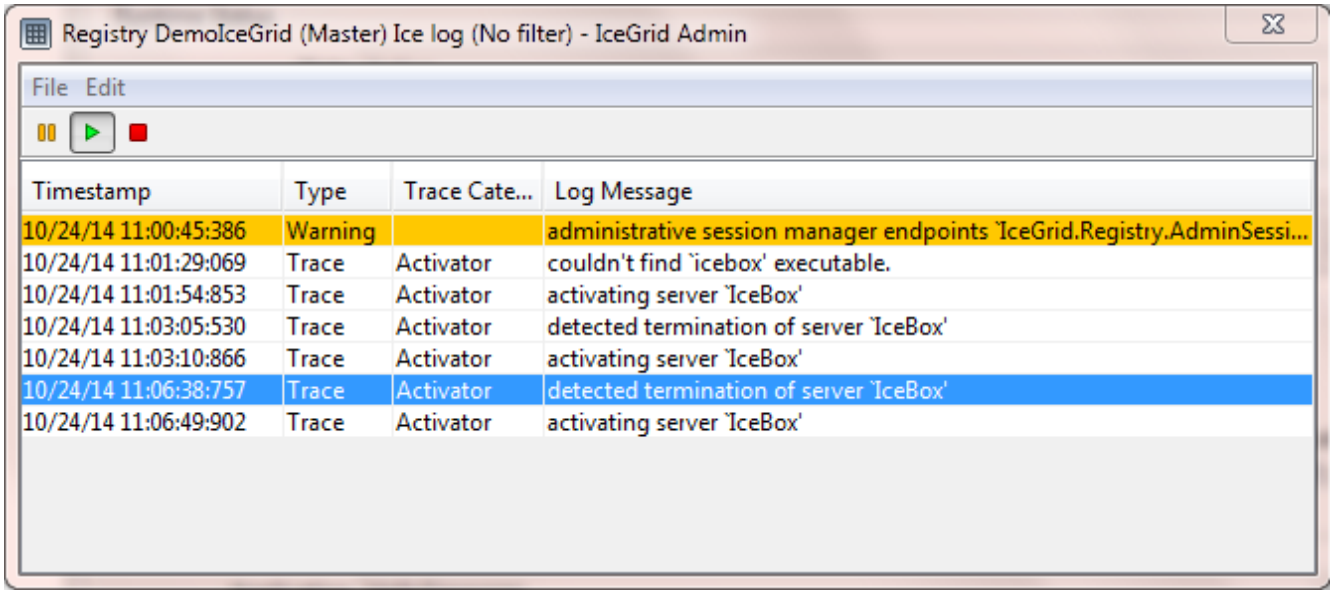

Garbage Collection Improvements

Ice for C++ has included a garbage collection facility for many years, but it suffered from one significant flaw: it was impossible to guarantee that object graphs could be collected in a thread-safe manner. In this release, we've improved the implementation and modified the API to ensure it's efficient and thread-safe. The application is now responsible for telling the Ice run time when an object graph is eligible for collection by calling the `Object::ice_collectable(bool)` method. Ice assumes that a graph eligible for collection is immutable, therefore you must not update the structure of the graph after marking it as collectable. The code below demonstrates the new API:

C++

```
NodePtr root = new Node();
NodePtr child = new Node();
root->left = child;
child->left = root; // Create a cycle

// Mark the graph as collectable. At this point you should no longer
// modify the structure of the graph referenced by "root".
root->ice_collectable(true);
```

Users of previous Ice versions should refer to the upgrade guide for details on what has changed.

Transport Improvements

Aside from adding the WebSocket transports and overhauling our SSL plug-in, we've made a few more enhancements to the Ice transports.

HTTP Proxy

Outgoing TCP and SSL connections can now be mediated by an HTTP network proxy service that supports HTTP CONNECT tunneling. The new properties `Ice.HTTPProxyHost` and `Ice.HTTPProxyPort` configure the addressing information for the proxy.

Source Address

Proxy endpoints can now include the `--sourceAddress` option to specify a network interface on which to bind outgoing connections. The `Ice.Default.SourceAddress` property supplies a default value when the endpoint option is not used.

Copyright 2016, ZeroC, Inc.

Batch Invocations

We made some significant changes to the way batch invocations work in Ice 3.6, making them more reliable and more consistent with other invocation modes. Please review the upgrade guide for details.

Collocated Invocations

Collocated invocations have been reimplemented to provide location transparency semantics that are much more consistent with those of regular remote invocations. For example, you can now make asynchronous invocations on a collocated servant, and the Python language mapping now supports collocated invocations.

If your application relies on collocated invocations, you should review our upgrade guide for more information.

Finder Interfaces

Router and locator implementations are now required to implement a *finder* interface that allows a client to obtain a proxy for the service while knowing only its address and port. A finder object has a well-known identity that a client can combine with addressing information to compose its initial proxy. Using the finder proxy, the client can request a proxy for the target service. For example, the IceGrid administrative utility only needs to prompt the user for addressing details but does not need to ask for the identity of the service's well-known object, which may have been customized in the service's configuration.

In addition to routers and locators, IceStorm also supports a finder interface for obtaining the topic manager proxy.

Java Interrupts

Ice for Java now allows an application to safely interrupt a thread that's blocked in a call to the Ice run time. Attempting to use interrupts in previous Ice releases led to undefined and generally undesirable behavior, but in Ice 3.6 the semantics are clearly defined. Consider this example:

Java

```
HelloPrx helloProxy = ...;
try {
    helloProxy.sayHello();
} catch (Ice.OperationInterruptedException ex) {
    // interrupted by another thread!
}
```

Here we've shown how an application can deal with an interrupt that occurs during a remote invocation. Calls to other blocking Ice APIs, such as `Communicator.waitForShutdown()`, can also be interrupted. Consult the manual for a complete description of the semantics and a list of the valid interruption points.

Java Lambda Functions

Ice 3.6 adds support for Java 8, which means you can now use lambda functions for asynchronous callbacks. As an example of how lambda functions can simplify your code, consider these two equivalent invocations:

Copyright 2016, ZeroC, Inc.

Java

```
// Using an anonymous callback object
p.begin_ice_ping(new Ice.Callback_Object_ice_ping()
    {
        @Override
        public void response()
        {
            System.out.println("done!");
        }
        @Override
        public void exception(Ice.LocalException ex)
        {
            System.out.println("error!");
        }
    });

// Using Java 8 lambdas
p.begin_ice_ping(
    () -> System.out.println("done!"),
    (Ice.Exception ex) -> System.out.println("error!"));
```

Java Buffers

The new metadata tag `"java:buffer"` can be applied to sequences of certain primitive types, causing the translator to use subclasses of `java.nio.Buffer` instead of the default mapping to a native Java array. The buffer mapping is especially convenient when an application already has data in a buffer object because it eliminates the need to create and fill a temporary array just to call a Slice operation. The buffer mapping also avoids copying when receiving these sequences by directly referencing Ice's unmarshaling buffer.

Here's a simple example:

Slice

```
sequence<byte> Bytes;
sequence<int> Ints;

interface Sender {
    ["java:buffer"] Ints translate(["java:buffer"] Bytes input);
};
```

The `translate` operation maps to the following Java method:

Java

```
java.nio.IntBuffer translate(java.nio.ByteBuffer input);
```

On the client side, a call to the `translate` proxy method returns an `IntBuffer` view of the data in the results buffer. On the server side, the implementation receives a `ByteBuffer` wrapper of the relevant portion of the incoming parameter buffer.

Copyright 2016, ZeroC, Inc.

C++ View Types

The new `cpp:type:view-type` metadata directive allows you to safely use "view" C++ types for operation parameters. Such C++ types point to memory held by other objects and don't copy memory. A typical example is the experimental C++14 `string_view` type. `cpp:type :view-type` can be applied to string, sequence and dictionary parameters, to speed-up your application by reducing the number of copies made during remote invocations. See the cpp:type and cpp:view-type Metadata Directives for details.

Linking with C++ Libraries on Windows

You no longer need to list Ice import libraries (such as `IceD.lib`) when linking with Ice C++ libraries on Windows. `pragma` directives in Ice header files provide these import library names to the linker, for example:

C++

```
// For Ice C++ library
#if defined(_MSC_VER)
#   if !defined(ICE_BUILDING_ICE)
#      if defined(_DEBUG) && !defined(ICE_OS_WINRT)
#          pragma comment(lib, "IceD.lib")
#      else
#          pragma comment(lib, "Ice.lib")
#      endif
#   endif
#endif
```

The Ice Builder for Visual Studio automatically configures the `$(IceLib)` directory that contains these import libraries. `$(IceLib)` depends on the versions of Ice and Visual Studio you're using, and on the target architecture. For example, `$(IceLib)` for a Ice 3.6, Visual Studio 2012 and a x64 target is `$(IceHome)\lib\vc110\x64`.

C++ Plug-in for Crypt Password Files

Both the Glacier2 router and IceGrid registry provide a simple file-based authentication mechanism using a "crypt password" file that contains a list of user name and password-hash pairs.

In Ice 3.5 and earlier releases, the Glacier2 router and IceGrid registry use the archaic DES-based Unix Crypt algorithm to hash the provided password and verify if this hash matches the hash in the password-file. This DES Unix Crypt implementation was provided by OpenSSL on all platforms.

Ice 3.6 supports more recent and secure hash formats and algorithms:

- Windows and OS X: PBKDF2 with SHA-1, SHA-256, or SHA-512 as digest algorithm
- Linux: SHA-256 and SHA-512 Crypt, plus the old DES-based Unix Crypt for compatibility with earlier releases

Windows and OS X no longer rely on OpenSSL for this password hashing, nor for anything else.

In Ice 3.6, we have also moved the implementation of this authentication mechanism to a separate C++ plug-in, the `Glacier2CryptPermissionsVerifier` plug-in. This plug-in is loaded automatically by Glacier2 or IceGrid if you configure a crypt passwords file using one of the existing `CryptPasswords` properties `Glacier2.CryptPasswords`, `IceGrid.Registry.AdminCryptPasswords` and `IceGrid.Registry.CryptPasswords`. If you don't use the simple authentication mechanism provided by the `CryptPasswords` properties, you don't need to include this plug-in with your application.

Support for -fvisibility=hidden with GCC and clang

Ice for C++ is now built by default with `-fvisibility=hidden` on Linux (with GCC) and OS X (with clang). You can also build code generated by `slice2cpp` with `-fvisibility=hidden`, and use the `--dll-export` option to export symbols from such generated code.

Copyright 2016, ZeroC, Inc.

Backward Compatibility of Ice Versions

The subsections below describe how Ice maintains (or does not maintain) backwards compatibility between versions.

On this page:

- Source-code compatibility
- Binary compatibility
- On-the-wire compatibility
- Database compatibility
- Interface compatibility
  - IceGrid
  - IceStorm

Source-code compatibility

Ice maintains source-code compatibility between a patch release (e.g., 3.6.3) and the most recent minor release (e.g., 3.6.0), but does not guarantee source-code compatibility between minor releases (e.g., between 3.5 and 3.6).

Upgrading your Application from Ice 3.5 describes the significant API changes in this release that may impact source-code compatibility. Furthermore, the subsections Removed APIs and Deprecated APIs summarize additional changes to Ice APIs that could affect your application.

Binary compatibility

As for source-code compatibility, Ice maintains backward binary compatibility between a patch release and the most recent minor release, but does not guarantee binary compatibility between minor releases.

The requirements for upgrading to a new minor (or major) release depend on the language mapping used by your application:

- For statically-typed languages (C++, Java, .NET), the application must be recompiled.
- For scripting languages that use static translation, your Slice files must be recompiled.
- No action is necessary for a Python or Ruby script that loads its Slice files dynamically.

On-the-wire compatibility

Ice always maintains "on the wire" compatibility with prior releases. A client using Ice version *x* can communicate with a server using Ice version *y* and vice versa.

The Ice message format is governed by two sets of rules: the protocol rules and the encoding rules. The protocol rules define the number of message types as well as the format of the message headers. The encoding rules specify how Slice types are marshaled. Ice maintains separate versions for the protocol and encoding, which makes it possible for them to evolve independently.

Several features introduced in Ice 3.5 required changes to the way that Slice classes and exceptions are marshaled. These changes make the encoding incompatible with previous releases, and therefore Ice 3.5 introduced version 1.1 of the Ice encoding. No additional encoding changes were made in Ice 3.6.

The protocol remains the same since the first version Ice, with protocol version number 1.0.

Changing the protocol or encoding format can be disruptive and does not happen often. Naturally, these changes raise the issue of backward compatibility with applications that use earlier versions of Ice, and more specifically, applications that use version 1.0 of the Ice encoding. Ice 3.5 and 3.6 use the following semantics:

- Encoding version 1.1 is the default for these releases. You can change this default using the `Ice.Default.EncodingVersion` property.
- A proxy is marshaled with an embedded encoding version representing the highest version supported by the target object, so the receiver of a proxy always knows the encoding version(s) it is allowed to use when invoking operations on that proxy. A proxy created locally uses the default encoding version, which can be inspected and changed using proxy methods.
- The results of an operation are encoded using the same version as the incoming request.
- As you would expect, the features provided by the 1.1 encoding are unavailable when an Ice application is forced to use the 1.0 encoding.

See Encoding Version 1.1 for additional details.

Copyright 2016, ZeroC, Inc.

Database compatibility

Upgrading to a new minor release of Ice often includes an upgrade to the supported version of Berkeley DB. In turn, this may require an application to migrate its databases, either because the format of Berkeley DB's database files has changed, or due to a change in the schema of the data stored in those databases.

For example, if your application uses Freeze, it may be necessary for you to migrate your databases even if your schema has not changed.

Certain Ice services also use Freeze in their implementation. If your application uses these services (IceGrid and IceStorm), it may be necessary for you to migrate their databases as well.

Please refer to the relevant subsections below for migration instructions.

Interface compatibility

Although Ice always maintains compatibility at the protocol level, changing Slice definitions can also lead to incompatibilities. As a result, Ice maintains interface compatibility between a patch release and the most recent minor release, but does not guarantee compatibility between minor releases.

This issue is particularly relevant if your application uses Ice services such as IceGrid or IceStorm, as a change to an interface in one of these services may adversely affect your application.

Interface changes in an Ice service can also impact compatibility with its administrative tools, which means it may not be possible to administer an Ice 3.6.x service using a tool from a previous minor release (or vice-versa).

IceGrid

Starting with Ice 3.2.0, IceGrid registries and nodes are interface-compatible so you can use different IceGrid registry and node versions within the same deployment. For example, you can use an IceGrid node from Ice 3.2 with a registry from Ice 3.4.

IceGrid master and slaves can have different versions as long as they use the same database format. Since the IceGrid database format changed in Ice 3.6, an IceGrid slave from earlier versions will not be able to synchronize with an IceGrid master from Ice 3.6 or later. You need to upgrade all your IceGrid registries to Ice 3.6 to ensure replication between the registries work.

An IceGrid node can activate Ice servers that use different Ice releases, for example, an IceGrid node version 3.6 can activate an Ice 3.4 server.

The IceGrid graphical and command-line administrative tools can only administer an IceGrid registry with the same `major.minor` version. For example, you have to use the 3.6 administrative tools to administer an IceGrid 3.6 registry; you cannot use an administrative tool from an earlier Ice version. The reverse is also true: you cannot administer an IceGrid 3.5 registry with an IceGrid 3.6 administration tool.

IceStorm

Topic linking is supported between all IceStorm versions released after 3.0.0.

Copyright 2016, ZeroC, Inc.

Upgrading your Application from Ice 3.6.0 or later

The subsections below provide information about upgrading existing Ice 3.6-based applications to Ice 3.6.3.

On this page:

- Full Source and Binary Compatibility
- Policy Assemblies for .NET Applications

Full Source and Binary Compatibility

Ice 3.6.3 maintains full source and binary compatibility with Ice 3.6.2, Ice 3.6.1 and Ice 3.6.0: you can upgrade your application from Ice 3.6.0 or later to Ice 3.6.3 without recompiling or relinking anything.

The database formats used by Ice services such as IceGrid and IceStorm have not changed, therefore no database migration is required. Generally speaking, you are free to use any combination of Ice 3.6.0, Ice 3.6.1, Ice 3.6.2 or Ice 3.6.3 applications and Ice services.

Some internal Ice interfaces may change in a patch release such as Ice 3.6.3, and therefore you must always upgrade all the Ice components of a given application together. For example, you should not use version 3.6.3 of the `ice36.dll` with version 3.6.0 of the `icessl36.dll` in the same application.

Policy Assemblies for .NET Applications

For a .NET application, Ice for .NET includes policy assemblies that supply the .NET run time with the required compatibility information. Policy assemblies have names of the form `policy.3.6.`*`package`*`.dll`. For example, the policy assembly for IceBox is `policy.3.6.IceBox.dll`. One way to upgrade an existing .NET application to a new patch release while maintaining binary compatibility is to install the policy assemblies into the Global Assembly Cache (GAC) with `gacutil`:

```
gacutil -i <policy.dll>
```

Another option is to modify the `.config` file of your application to add `bindingRedirect` directives, as explained in the links below:

http://msdn.microsoft.com/en-us/library/7wd6ex19.aspx
http://msdn.microsoft.com/en-us/library/yx7xezcf.aspx

For example, in the `.config` file of an application you can modify the configuration of the Ice and IceBox assemblies as follows:

Copyright 2016, ZeroC, Inc.

XML

```
<?xml version="1.0" encoding="utf-8" ?>
<configuration>
<runtime>
<assemblyBinding xmlns="urn:schemas-microsoft-com:asm.v1">

<dependentAssembly>
  <assemblyIdentity name="Ice" culture="neutral"
publicKeyToken="cdd571ade22f2f16"/>
  <bindingRedirect oldVersion="3.6.0.0" newVersion="3.6.3.0"/>
</dependentAssembly>

<dependentAssembly>
  <assemblyIdentity name="Ice" culture="neutral"
publicKeyToken="cdd571ade22f2f16"/>
  <bindingRedirect oldVersion="3.6.1.0" newVersion="3.6.3.0"/>
</dependentAssembly>

<dependentAssembly>
  <assemblyIdentity name="Ice" culture="neutral"
publicKeyToken="cdd571ade22f2f16"/>
  <bindingRedirect oldVersion="3.6.2.0" newVersion="3.6.3.0"/>
</dependentAssembly>

<dependentAssembly>
  <assemblyIdentity name="IceBox" culture="neutral"
publicKeyToken="cdd571ade22f2f16"/>
  <bindingRedirect oldVersion="3.6.0.0" newVersion="3.6.3.0"/>
</dependentAssembly>

<dependentAssembly>
  <assemblyIdentity name="IceBox" culture="neutral"
publicKeyToken="cdd571ade22f2f16"/>
  <bindingRedirect oldVersion="3.6.1.0" newVersion="3.6.3.0"/>
</dependentAssembly>

<dependentAssembly>
  <assemblyIdentity name="IceBox" culture="neutral"
publicKeyToken="cdd571ade22f2f16"/>
  <bindingRedirect oldVersion="3.6.2.0" newVersion="3.6.3.0"/>
</dependentAssembly>

</assemblyBinding>
</runtime>
</configuration>
```

This `.config` file can be used for `iceboxnet.exe` (as `iceboxnet.exe.config`), to load a simple IceBox service built with Ice 3.6.0 with `iceboxnet` 3.6.3.

Note that `cdd571ade22f2f16` is the token corresponding to ZeroC's public key for signing the assemblies in binary distributions. If you build Ice from sources, your assemblies are signed using the development key found in `config/IceDevKey.snk`. The token for this development key is `1f998c50fec78381`.

Copyright 2016, ZeroC, Inc.

The advantage of installing the policy assemblies into the GAC is that they establish binary compatibility for all Ice applications, whereas modifying a `.config` file must be done for each application individually.

Copyright 2016, ZeroC, Inc.

Upgrading your Application from Ice 3.5

The subsections below provide additional information about upgrading to Ice 3.6, including administrative procedures for the supported platforms.

On this page:

- Timeout changes
- ACM changes
- IceGrid and Glacier2 sessions
- Glacier2 compatibility for helper classes
- SSL changes
  - SSLv3 disabled by default
  - Certificate verification
  - SSL changes on OS X
  - SSL changes on Windows
  - SSL changes on Linux
- Collocated Invocation changes
- Batch Invocation changes
- Logger changes
- Crypt Password changes
- C++ changes
  - Garbage collection changes
  - String converter changes
  - OS X with C++
- Java changes
  - Java mapping changes
  - New Java features
- C# changes
  - C# mapping changes
  - C# serialization changes
- Python changes
- PHP changes
- Ruby changes
- JavaScript changes
  - JavaScript mapping changes
  - JavaScript packaging changes
- Objective-C changes
- Ubuntu packages
- Migrating IceGrid databases from Ice 3.5
- Migrating IceStorm databases from Ice 3.5
- Migrating Freeze databases from Ice 3.5
- Migrating Android applications from Ice 3.5
- Changed APIs
- Removed APIs
- Deprecated APIs
- Visual C++ compiler warnings

Timeout changes

In previous Ice versions, the timeout set for a connection also served as the timeout for all invocations on that connection, such that an invocation timeout would cause Ice to close the connection and consequently report a timeout exception for all pending invocations on that connection.

With the addition of invocation timeouts, Ice now provides a much cleaner separation between two distinct features:

- Connection timeouts
  These timeouts should be used as a fail-safe strategy for handling unrecoverable network errors in a timely fashion.

- Invocation timeouts
  You can now safely abort an invocation that's taking too long to complete without affecting other invocations pending on the same connection.

Together with the new heartbeat functionality offered by Ice's Active Connection Management facility, applications now have much more control over their connections.

Copyright 2016, ZeroC, Inc.

Existing Ice applications could configure connection timeouts in several ways, such as by setting Ice.Override.Timeout, or with the `-t` option in endpoints, or by calling `ice_timeout` on proxies. Ice still supports all of these options, and they all configure connection timeouts just like in previous Ice versions. However, you should carefully review these settings to see that they match the true purpose of connection timeouts. Specifically, connection timeouts should normally be chosen based on the speed of the network on which a connection takes place; a small timeout value is fine for a relatively fast network, but a larger value may be necessary for slower connections. Ice enforces connection timeouts when performing network operations such as connection establishment, reading, writing, and closing; setting a timeout allows Ice (and therefore the application) to detect a network problem within a well-defined time period without waiting for low-level network protocols to detect and report the issue.

We recommend that every application enable connection timeouts, and as a result they are now enabled by default as defined by the new property `Ice.Default.Timeout`.

The timeout semantics in previous Ice versions made it difficult for developers to employ invocation timeouts correctly. Using a small timeout value in order to detect network issues reasonably soon risked the possibility that some invocations might time out prematurely. These conflicting goals may have caused developers to modify their application designs to avoid long-running invocations. With the introduction of a separate invocation timeout feature, it's now possible to set connection timeouts appropriately for network errors while using invocation timeouts only for those operations that require them. Unlike connection timeouts, where setting different timeout values causes Ice to open separate connections, invocation timeouts of various values can be freely mixed on the same connection.

The default invocation timeout in Ice 3.6 is "infinite", meaning invocations do not time out by default. You can change the default setting using the new property `Ice.Default.InvocationTimeout`. For individual proxies, the method `ice_invocationTimeout` returns a new proxy configured with the desired invocation timeout. If you need to continue using the timeout semantics of previous Ice versions, set `Ice.Default.InvocationTimeout` to `-2` or call `ice_invocationTimeout` with the value `-2` on a given proxy.

Note that retry semantics differ between connection and invocation timeouts. Ice still performs automatic retries if possible for connection timeouts, but does not retry invocations that fail due to an invocation timeout.

ACM changes

The Active Connection Management (ACM) facility now offers additional control over its behavior, along with a new automatic heartbeat feature. Client-side ACM was enabled by default in prior Ice versions, which means an existing Ice application most likely uses ACM unless the application explicitly disabled it.

There are several ACM changes that may affect an existing application:

- The `Ice.ACM.Client` and `Ice.ACM.Server` properties have been deprecated and replaced by `Ice.ACM.Client.Timeout` and `Ice.ACM.Server.Timeout`, respectively.
- With the new heartbeat feature, Ice automatically sends heartbeat messages at regular intervals. You can use this feature in place of a dedicated background thread that was commonly used in previous versions of Ice for keeping a session or a bidirectional connection active.
- It's no longer necessary to disable ACM in Glacier2 clients as long as you enable ACM heartbeats.
- Server-side ACM is now enabled by default. Like with the previous versions, the default server-side configuration doesn't close idle connections. However, it now enables heartbeats while incoming requests are pending. This ensures the client doesn't close the connection prematurely while there are long invocations pending.

IceGrid and Glacier2 sessions

The ACM changes in Ice 3.6 offer new implementation strategies for applications that create Glacier2 or IceGrid sessions. Since a session is tightly bound to a connection, we strongly recommended in prior releases that you disable ACM altogether to avoid the risk of ACM prematurely closing a connection and consequently terminating a session. As of Ice 3.6, it's no longer necessary to disable ACM in these situations. Furthermore, the addition of ACM heartbeats means you can remove existing code that creates a background thread just to keep a session alive.

The `Ice::Connection` interface provides several new operations, including the ability to obtain and modify a connection's current ACM settings. If you use the Glacier2 helper classes, they call the `Connection` operations to tailor the ACM timeout and heartbeat based on the router's ACM configuration. Applications that manually create a Glacier2 session can configure ACM like this:

Copyright 2016, ZeroC, Inc.

C++

```
Glacier2::RouterPrx router = ...
// Create a session...
int acmTimeout = router->getACMTimeout();
if(acmTimeout > 0)
{
    Ice::ConnectionPtr connection = router->ice_getCachedConnection();
    connection->setACM(acmTimeout, IceUtil::None, Ice::HeartbeatAlways);
}
```

The new operation `Glacier2::Router::getACMTimeout` returns the router's server-side ACM timeout. In this example, the client calls `setACM` on its connection to the router, passing this same timeout value to ensure consistency between client and server. The client also enables automatic heartbeats so that the connection remains active and prevents the router's server-side ACM from closing the connection.

The ACM improvements include changes to the ACM configuration properties. You can use these properties to achieve the same result as the code above, however the properties can potentially affect other connections as well.

Finally, another new operation on the `Connection` interface lets you specify a callback object that will be notified when the connection receives a heartbeat message, and when the connection closes. This feature can be especially useful for session-based applications that need to closely monitor their connections.

Glacier2 compatibility for helper classes

The Glacier2 helper classes in Ice 3.6 now depend on the ACM heartbeat features described above to keep a session alive. This functionality requires a Glacier2 router that also uses Ice 3.6 or later. If you're upgrading a client to Ice 3.6, we strongly recommend upgrading the Glacier2 router to 3.6 as well. Using an earlier version of Glacier2 will require your application to manually keep the session alive.

SSL changes

IceSSL for C++ has been overhauled to make use of platform-native SSL APIs where possible:

- IceSSL on Windows now uses SChannel
- IceSSL on OS X now uses Secure Transport
- Linux platforms continue to use OpenSSL as in previous releases

As a result, there have been a number of changes to the IceSSL configuration properties and its C++ API. In the IceSSL property reference, you'll see platform differences marked with **C++ using Windows**, **C++ using OS X**, and **C++ using OpenSSL**, respectively.

We discuss the affected platforms below, along with other general SSL changes.

These changes also affect applications that use Ice for Python, Ice for Ruby, and Ice for PHP because these language mappings are based on Ice for C++.

SSLv3 disabled by default

To improve security, IceSSL now disables the SSLv3 protocol by default. In other words, only TLS protocols are enabled by default.

Although we do not recommend it, you can enable SSLv3 using the following settings:

Copyright 2016, ZeroC, Inc.

```
# Enable only SSLv3
IceSSL.Protocols=SSL3

# Enable SSLv3 and TLS
IceSSL.Protocols=SSL3, TLS1_0, TLS1_1, TLS1_2

# OS X only: Enables SSLv3 and TLS
IceSSL.ProtocolVersionMin = "SSLv3"
```

Refer to IceSSL.* for more information on these settings.

Certificate verification

The default value of the `IceSSL.VerifyDepthMax` property is now three (it was previously two). This allows certificate chains of three certificates, such as a chain consisting of a peer certificate, a CA certificate, and a Root CA certificate.

The certificate chain provided in `IceSSL::ConnectionInfo` should now always include the root certificate if the chain is successfully verified.

We also added a `verified` member to `IceSSL::ConnectionInfo`, which indicates whether or not the peer certificate was successfully verified. For a client that sets `IceSSL.VerifyPeer=0`, this member allows the client to check the verification status of the server's certificate. For server connections, the member should always be true since servers always reject invalid client certificates.

SSL changes on OS X

By using Secure Transport, IceSSL on OS X has now become very similar to IceSSL in Ice Touch. For example, you can use OS X keychains, take advantage of the system's default graphical password prompt, and use many of the same configuration properties.

General Changes

- Password prompt
  In previous releases, OpenSSL would attempt to prompt command-line users for a password if the application failed to define a password or supply a password callback. Now OS X will use its default graphical password prompt in this situation.
- DSA certificates
  DSA certificates are no longer supported on OS X. If you used DSA certificates, you will need to generate RSA equivalents. The `IceSSL.CertFile` property still accepts the same syntax in that it allows you to specify two files (one RSA certificate file and one DSA certificate file), but it will ignore the DSA certificate.
- MD5 signatures
  OS X does not support certificates with MD5 signatures. We recommend using SHA256 instead.

API Changes

- Since IceSSL on OS X no longer uses OpenSSL, the native C++ class `IceSSL::Plugin` does not support the `setContext` and `getContext` methods.
- The method `IceSSL::Certificate::verify(const PublicKeyPtr&)` has been deprecated. The new method `IceSSL::Certificate::verify(const CertificatePtr&)` takes its place.

Property Changes

- Keychains
  The properties `IceSSL.Keychain` and `IceSSL.KeychainPassword` are now supported on OS X.
- Certificate authorities
  Use the new property `IceSSL.CAs` to denote a file containing the certificates of your trusted certificate authorities. If you'd rather use the system's default certificate authorities, enable `IceSSL.UsePlatformCAs` instead.
- Certificates
  Use the `IceSSL.CertFile` property to denote a PKCS12 file containing both a certificate and a private key. The `IceSSL.KeyFil`

Copyright 2016, ZeroC, Inc.

e property is now deprecated and should no longer be used to denote a separate private key file.

- Certificate queries
  The new property `IceSSL.FindCert` lets you query a keychain to find a certificate matching certain criteria.
- Diffie Hellman parameters
  The new property `IceSSL.DHParams` allows you to specify the name of a file containing pre-generated Diffie Hellman parameters. The property `IceSSL.DH.bits` is no longer supported on OS X.
- Protocol limits
  The `IceSSL.Protocols` property is no longer supported on OS X for specifying the versions of SSL/TLS that a program will accept. Use the new properties `IceSSL.ProtocolVersionMin` and `IceSSL.ProtocolVersionMax` instead.

Refer to deprecated APIs for information about obsolete properties.

SSL changes on Windows

By using SChannel, IceSSL on Windows has now become very similar to IceSSL for .NET. For example, you can use certificate stores and many of the same configuration properties.

General Changes

- Password prompt
  In previous releases, OpenSSL would attempt to prompt command-line users for a password if the application failed to define a password or supply a password callback. Windows does not have a default password prompt, so this situation will now result in a run-time exception.
- Anonymous Diffie Hellman (ADH)
  ADH ciphers are no longer supported on Windows.
- Certificate formats
  Windows is able to load a certificate in PEM format if no password is required. Use the PFX (PKCS#12) format for password-protected certificates and keys.

API Changes

- Since IceSSL on Windows no longer uses OpenSSL, the native C++ class `IceSSL::Plugin` does not support the `setContext` and `getContext` methods.
- The method `IceSSL::Certificate::verify(const PublicKeyPtr&)` has been deprecated. The new method `IceSSL::Certificate::verify(const CertificatePtr&)` takes its place.

Property Changes

- Certificate authorities
  Use the new property `IceSSL.CAs` to denote a file containing the certificates of your trusted certificate authorities. If you'd rather use the system's default certificate authorities, enable `IceSSL.UsePlatformCAs` instead.
- Certificates
  Use the `IceSSL.CertFile` property to denote a PKCS12 file containing both a certificate and a private key. The `IceSSL.KeyFile` property is now deprecated and should no longer be used to denote a separate private key file.
- Certificate queries
  The new property `IceSSL.FindCert` lets you query a certificate store to find a certificate matching certain criteria.

Refer to deprecated APIs for information about obsolete properties.

SSL changes on Linux

API Changes

- The method `IceSSL::Certificate::verify(const PublicKeyPtr&)` has been deprecated. The new method `IceSSL::Certificate::verify(const CertificatePtr&)` takes its place.

Property Changes

Copyright 2016, ZeroC, Inc.

- Certificate authorities
  Use the new property `IceSSL.CAs` to denote a file containing the certificates of your trusted certificate authorities. If you'd rather use the system's default certificate authorities, enable `IceSSL.UsePlatformCAs` instead.

- Certificates
  Use the `IceSSL.CertFile` property to denote a PKCS12 file containing both a certificate and a private key. The `IceSSL.KeyFile` property is now deprecated and should no longer be used to denote a separate private key file.

Refer to deprecated APIs for information about obsolete properties.

Collocated Invocation changes

Ice has always supported location transparency, where a proxy invocation should ideally have the same semantics regardless of whether the target object is on a different host, a different process on the same host, or collocated in the current process. For this latter case, a collocated invocation is defined as a proxy invocation on a target object that is hosted by an object adapter in the same process and created using the same communicator as the proxy. In previous versions of Ice, the semantics of collocated invocations differed in several ways from regular "remote" invocations:

- Classes and exceptions were never sliced. Instead, the receiver always received a class or exception as the derived type that was sent by the sender.
- If a collocated invocation threw an exception that was not in an operation's exception specification, the original exception was raised in the client instead of `UnknownUserException`. (This applied to the C++ mapping only.)
- Class factories were ignored.
- Invocation timeouts were ignored.
- If an operation implementation used an in parameter that was passed by reference as a temporary variable, the change affected the value of the in parameter in the caller (instead of modifying a temporary copy of the parameter on the callee side only).
- Asynchronous semantics were not supported for collocated invocations.

Ice 3.6 eliminates all of these differences. Most notably, you can now use the asynchronous invocation API on a collocated servant, and collocated invocations are now supported in Python.

Note however that a few differences in semantics still remain:

- Most collocated invocations are dispatched in the server-side thread pool just like regular invocations; the only exceptions are synchronous twoway collocated invocations with no invocation timeout, which are dispatched in the calling thread.
- The state of the servant's object adapter is ignored: collocated invocations proceed normally even if the servant's adapter is not yet activated or is in the holding state.
- AMI callbacks for asynchronous collocated invocations are dispatched from the servant's thread and not from the client-side thread pool unless AMD is used for the servant dispatch. In this case, the AMI callback is called from the client-side thread pool.
- Invocation timeouts work as usual, but connection timeouts are ignored.

If your application relies on collocated invocations, test it carefully with Ice 3.6 to ensure that it still behaves as expected.

Batch Invocation changes

We made some significant changes to the way batch invocations work in Ice 3.6. The following table compares the behavior of batch invocations between previous releases and Ice 3.6:

Batch invocation behavior in Ice 3.5 and earlier	Batch invocation behavior in Ice 3.6
A batch request is queued by the connection associated with the proxy.	A batch request is queued by the proxy used for the invocation. The only exception is for a fixed proxy , in which case batch requests continue to be queued by the connection associated with the proxy.
If a proxy was not already associated with a connection, the initial batch request could trigger a connection attempt.	A batch request does not trigger any network activity.
Calling `Communicator::flushBatchRequests` flushes all queued requests for all connections.	Calling `Communicator::flushBatchRequests` only flushes queued requests invoked using fixed proxies.

Copyright 2016, ZeroC, Inc.

Calling `Connection::flushBatchRequests` flushes queued requests for all proxies associated with that connection.	Calling `Connection::flushBatchRequests` only flushes queued requests invoked using fixed proxies associated with that connection.
A batch of requests is compressed when at least one of the proxies used to create the requests in this batch has compression enabled.	A batch of requests is compressed when the proxy used to flush this batch has compression enabled. Batched requests flushed through `Communicator::flushBatchRequests` or `Connection::flushBatchRequest` never use compression.
Batch requests are not affected by proxy lifecycles because the requests are queued by connections.	Batch requests queued by a proxy are discarded when the proxy is deallocated.
Batch requests could be silently lost if an error occurred during a flush.	Calling `ice_flushBatchRequests` on a regular (non-fixed) proxy behaves like a oneway request: failures that occur during network activity trigger automatic retries and can eventually raise an exception. Furthermore, the `sent` callback is invoked for asynchronous calls to `ice_flushBatchRequests`.

With these changes, Ice's batch invocation facility has become more reliable and behaves more consistently with other invocation modes.

To give you more control over batch requests, we've also added a new `BatchInterceptor` API. You can configure a communicator with a custom batch interceptor in order to implement your own auto-flush algorithms or to receive notification when an auto-flush fails.

Logger changes

We added a new operation `getPrefix` to the local interface `Logger`.

Slice

```
module Ice {
local interface Logger {
...
     string getPrefix();
};
};
```

If you implement your own logger, you will need to update your implementation with a new `getPrefix`. `getPrefix` returns the prefix associated with this logger.

Crypt Password changes

Both the Glacier2 router and IceGrid registry provide a simple file-based authentication mechanism using a "crypt password" file that contains a list of user name and password-hash pairs. In Ice 3.5 and earlier releases, the Glacier2 router and IceGrid registry use the archaic DES-based Unix Crypt algorithm to hash the provided password and verify if this hash matches the hash in the password-file.

Ice 3.6 no longer supports this hash format on Windows and OS X. As a result, you need to regenerate your crypt password files on these platforms when upgrading to Ice 3.6.

C++ changes

Garbage collection changes

We've made significant changes to the garbage collection facility for cyclic object graphs. If your application uses this facility, note that we've removed the following:

Copyright 2016, ZeroC, Inc.

- `Ice.GC.Interval` property
- `Ice.Trace.GC` property
- `Ice::collectGarbage()` function

The new property `Ice.CollectObjects` determines whether garbage collection is enabled by default for Slice class instances that are unmarshaled by the Ice run time. Refer to the garbage collection discussion for more information on using this feature.

String converter changes

A number of changes have been made to the C++ string conversion API in this release:

- The `stringConverter` and `wstringConverter` members of `Ice::InitializationData` have been removed. Applications should use the process-wide string converter API instead.
- The classes and functions have moved from the `Ice` namespace to the `IceUtil` namespace.
- Overloaded versions of the `nativeToUTF8` and `UTF8ToNative` convenience functions that accepted a communicator argument have been removed.
- The arguments have changed for the `stringToWstring` and `wstringToString` convenience functions.

OS X with C++

C++ developers on OS X need to be aware of several changes:

- C++11 libraries
  With Ice 3.5, C++11 applications needed to link with a separate set of C++11-specific Ice libraries located in `<Ice installation directory>/lib/c++11`. The Ice 3.6 binary distribution includes a single set of libraries that also support C++11, so you'll need to modify your application's library path to use `<Ice installation directory>/lib` instead.

- Minimum required version
  Ice 3.6 supports OS X 10.9 and 10.10, therefore the C++ libraries in the Ice binary distribution are built with `macosx-min-version=10.9`. Consequently, these libraries require `libc++` and are not compatible with `libstdc++`.

Java changes

Java mapping changes

The default constructor generated for Slice structures, exceptions and classes behaves differently for Ice 3.6 than in previous releases. Specifically, the default constructor now initializes string data members to an empty string, enumerator data members to the first enumerator in the enumeration, and structure data members to a default-constructed value.

In previous releases, the default constructor initialized these data members to null. Applications that depend on this behavior will require updating.

For situations where the overhead of default-constructing structure data members is undesirable, applications can call the one-shot constructor instead.

The Java mapping has also relaxed its marshaling prerequisites: it is no longer necessary to initialize data members of type structure or enumerator prior to an invocation. In Ice 3.6, passing a null value for a structure or enumerator causes Ice to marshal a default-constructed structure or the first enumerator, respectively.

New Java features

We have added several features that you may wish to incorporate into your Java application:

- Interrupts
- Buffers
- Java 8 and lambdas

C# changes

C# mapping changes

Copyright 2016, ZeroC, Inc.

The default constructor generated for Slice structures, exceptions and classes behaves differently for Ice 3.6 than in previous releases. Specifically, the default constructor now initializes string data members to an empty string and structure data members to a default-constructed value.

In previous releases, the default constructor did not explicitly initialize these data members and so they had the default C# values. Applications that depend on this behavior will require updating.

For situations where the overhead of default-constructing structure data members is undesirable, applications can call the one-shot constructor instead.

The C# mapping has also relaxed its marshaling prerequisites: it is no longer necessary to initialize data members of type structure prior to an invocation. In Ice 3.6, passing a null value for a structure causes Ice to marshal a default-constructed structure.

C# serialization changes

One of the new features in Ice 3.6 is support for .NET serialization for all Slice types (except proxies). Existing applications may be affected by this change because the `Ice.Optional` type now implements `ISerializable` and the serialized format of an optional value is different than in Ice 3.5.

Python changes

Developers who are migrating existing Ice applications from Ice 3.5 should be aware of a change that affects the Python language mapping. The `Ice.Unset` value now has `False` semantics, making it more convenient to test whether an optional Slice data member has a value:

Python

```
# Only valid with Ice 3.6!
if obj.optionalMember:
    # optionalMember has a value
```

With Ice 3.5, the code above would not have the intended behavior because the test would be true even when `optionalMember` is set to `Ice.Unset`. The correct way to write this using Ice 3.5 is shown below:

Python

```
# Correct test with Ice 3.5
if obj.optionalMember is not Ice.Unset:
    # optionalMember has a value
```

This code will also have the correct behavior with Ice 3.6, but the new style is easier to read. Also note that the Ice 3.6 semantics mean you need to use caution for optional values that can legally be set to `None`:

Python

```
if obj.optionalMember: # Fails for None AND Ice.Unset!
    # optionalMember is not Ice.Unset or None
```

You can distinguish between `Ice.Unset` and `None` as follows:

Copyright 2016, ZeroC, Inc.

Python

```
if obj.optionalMember is Ice.Unset:
    # optionalMember is unset
elif obj.optionalMember is None:
    # optionalMember is set to None
else:
    # optionalMember is set to a value other than None
```

We recommend that you review for correctness all uses of `Ice.Unset` and tests of optional data members.

PHP changes

No changes to the PHP mapping in Ice 3.6 affect compatibility for existing applications.

Ruby changes

No changes to the Ruby mapping in Ice 3.6 affect compatibility for existing applications.

JavaScript changes

JavaScript mapping changes

The default constructor generated for Slice structures, exceptions and classes behaves differently for Ice 3.6 than in previous releases. Specifically, the default constructor now initializes string data members to an empty string and structure data members to a default-constructed value.

In previous releases, the default constructor initialized these data members to null. Applications that depend on this behavior will require updating.

For situations where the overhead of default-constructing structure data members is undesirable, applications can call the one-shot constructor instead.

The JavaScript mapping has also relaxed its marshaling prerequisites: it is no longer necessary to initialize data members of type structure or enumerator prior to an invocation. In Ice 3.6, passing a null value for a structure or enumerator causes Ice to marshal a default-constructed structure or the first enumerator, respectively.

JavaScript packaging changes

The NodeJS packaging has changed from the original Ice for JavaScript 0.1 release, meaning existing JavaScript applications will need to modify their `require` statements. All of the top-level Ice modules (Ice, Glacier2, etc.) are now accessible by including the `icejs` package:

JavaScript

```
var Ice = require("icejs").Ice;
var Glacier2 = require("icejs").Glacier2;
// ...
```

Loading the generated code for your own Slice definitions looks similar. Suppose we have the following definitions in `Hello.ice`:

Copyright 2016, ZeroC, Inc.

Slice

```
module Demo {
interface Hello {
    idempotent void sayHello(int delay);
    void shutdown();
};
};
```

To make the `Demo` module conveniently accessible in our code, we can write:

JavaScript

```
var Demo = require("Hello").Demo;
var proxy = Demo.HelloPrx.uncheckedCast(...);
```

Nothing has changed for browser-based JavaScript applications, where loading `Ice.js` adds the Ice definitions to the global window object, and other Ice modules (`Glacier2.js`, etc.) must be loaded individually.

Note that Ice 3.6 adds support for the WebSocket transport to the Ice core, and includes new implementations of the WebSocket transport in Java and C#. This means you no longer need to use Glacier2 as an intermediary if your JavaScript client needs to communicate with a Java or C# server.

Objective-C changes

The default `init` method and convenience constructor generated for Slice structures, exceptions and classes behave differently for Ice 3.6 than in previous releases. Specifically, these methods now initialize string data members to an empty string, enumerator data members to the first enumerator, and structure data members to a default-constructed value.

In previous releases, the default constructor zero-initialized these data members. Applications that depend on this behavior will require updating.

For situations where the overhead of default-constructing structure data members is undesirable, applications can call the one-shot constructor instead.

The Objective-C mapping has also relaxed its marshaling prerequisites: it is no longer necessary to initialize data members of type structure prior to an invocation. In Ice 3.6, passing a null value for a structure causes Ice to marshal a default-constructed structure.

Ubuntu packages

Users of Ice 3.5 on Ubuntu had the choice of using Debian's packages or ZeroC's own experimental packages. We called them "experimental" because we expected we might eventually change the packaging structure, and in fact we have changed the structure for Ice 3.6.

Upgrading an existing installation of the ZeroC packages for Ice 3.5 on Ubuntu to Ice 3.6 is relatively straightforward. First add the Ice repository to the system and update the package list:

```
$ sudo apt-add-repository "deb https://zeroc.com/download/apt/ice/ubuntu14 stable main"
```

```
$ sudo apt-get update
```

To upgrade all of the run-time packages:

```
$ sudo apt-get install zeroc-ice-all-runtime
```

To upgrade all of the development packages:

Copyright 2016, ZeroC, Inc.

```
$ sudo apt-get install libzeroc-ice-dev libzeroc-ice-java zeroc-ice-all-dev
```

Refer to our binary distribution page for details on the individual packages.

Migrating IceGrid databases from Ice 3.5

Ice 3.6 supports the migration of IceGrid databases from Ice 3.3, 3.4 and 3.5. To migrate from earlier Ice versions, you will first need to migrate the databases to the Ice 3.3 format. If you require assistance with such migration, please contact support@zeroc.com.

To migrate, first stop the IceGrid registry you wish to upgrade.

Next, copy the IceGrid database environment to a second location:

```
$ cp -r db recovered.db
```

Locate the correct version of the Berkeley DB recovery tool (usually named `db_recover`). It is essential that you use the `db_recover` executable that matches the Berkeley DB version of your existing Ice release. For Ice 3.3, use `db_recover` from Berkeley DB 4.6. For Ice 3.4, use `db_recover` from Berkeley DB 4.8. For Ice 3.5, use `db_recover` from Berkeley DB 5.3. You can verify the version of your `db_recover` tool by running it with the `-V` option:

```
$ db_recover -V
```

Now run the utility on your copy of the database environment:

```
$ db_recover -h recovered.db
```

Change to the location where you will store the database environments for IceGrid 3.6:

```
$ cd <new-location>
```

Next, run the `upgradeicegrid36.py` utility located in the `config` directory of your Ice distribution (or in `/usr/share/Ice-3.6` if using an RPM installation). The first argument is the path to the old database environment. The second argument is the path to the new database environment.

In this example we'll create a new directory `db` in which to store the migrated database environment:

```
$ mkdir db
$ upgradeicegrid36.py <path-to-recovered.db> db
```

Upon completion, the `db` directory contains the migrated IceGrid databases.

By default, the migration utility assumes that the servers deployed with IceGrid also use Ice 3.6. If your servers still use an older Ice version, you need to specify the `--server-version` command-line option when running `upgradeicegrid36.py`:

```
$ upgradeicegrid36.py --server-version 3.5.1 <path-to-recovered.db> db
```

The migration utility will set the `server` descriptor attribute `ice-version` to the specified version and the IceGrid registry will generate configuration files compatible with the given version.

If you are upgrading the master IceGrid registry in a replicated environment and the slaves are still running, you should first restart the master registry in read-only mode using the `--readonly` option, for example:

```
$ icegridregistry --Ice.Config=config.master --readonly
```

Next, you can connect to the master registry with `icegridadmin` or the IceGrid administrative GUI from Ice 3.6 to ensure that the database is correct. If everything looks fine, you can shutdown and restart the master registry without the `--readonly` option.

IceGrid slaves from Ice 3.3, 3.4 or 3.5 won't interoperate with the IceGrid 3.6 master. You can leave them running during the upgrade of the master to not interrupt your applications. Once the master upgrade is done, you should upgrade the IceGrid slaves to Ice 3.6 using the instructions above.

Migrating IceStorm databases from Ice 3.5

No changes were made to the database schema for IceStorm in this release. Furthermore, Ice 3.5 and Ice 3.6 use the same version of Berkeley DB (Berkeley DB 5.3.x). You can use IceStorm databases created with Ice 3.5 with Ice 3.6 without any transformation.

Migrating Freeze databases from Ice 3.5

Copyright 2016, ZeroC, Inc.

No changes were made that would affect the content of your Freeze databases. Furthermore, Ice 3.5 and Ice 3.6 use the same version of Berkeley DB (Berkeley DB 5.3.x). You can use Freeze databases created with Ice 3.5 with Ice 3.6 without any transformation.

Migrating Android applications from Ice 3.5

Our recommended development environment for Android applications is Android Studio. Refer to the *Using the Binary Distribution* page appropriate for your platform for instructions on configuring a project in Android Studio.

Changed APIs

This section describes APIs whose semantics have changed, potentially in ways that are incompatible with previous releases.

The following APIs were changed in Ice 3.6:

- String conversion in C++
  Several changes were made to the string conversion API in C++.
- IceSSL
  The native IceSSL APIs changed on some platforms.
- C++ garbage collection
  The `Ice::collectGarbage` function was removed and replaced by a new garbage collection facility.

Removed APIs

This section generally describes APIs that were deprecated in a previous release and have now been removed. Your application may no longer compile or operate successfully if it relies on one of these APIs.

The following APIs were removed in Ice 3.6:

- Deprecated API for asynchronous method invocations (AMI)
  This API, which uses proxy operations such as `sayHello_async`, was deprecated in Ice 3.4 and is no longer available. The new API should be used instead. Refer to the appropriate language mapping section for more information.
- `Stats` facility
  This functionality was deprecated in Ice 3.5 and is now provided by the Instrumentation facility and the Metrics facet.
- `Ice.GC.Interval`
  `Ice.Trace.GC`
  `Ice::collectGarbage()`
  Significant changes were made to the C++ garbage collection facility.
- `Ice.MonitorConnections`
  This setting is no longer necessary.
- `IceSSL::Plugin::setContext()`
  `IceSSL::Plugin::getContext()`
  These C++ methods are no longer available on Windows or OS X.
- `Ice::InitializationData::stringConverter`
  `Ice::InitializationData::wstringConverter`
  These C++ data members are no longer available. Use `IceUtil::setProcessStringConverter` and `IceUtil::setProcessWstringConverter` instead.
- `Ice::Router::addProxy()`
  `IceGrid::Admin::writeMessage()`
  `IceStorm::Topic::subscribe()`
  These Slice operations were deprecated in previous Ice releases.
- `IceUtil.Version`
  This Java class was deprecated in previous Ice releases. Use `Ice.Util.stringVersion()` and `Ice.Util.intVersion()` instead.

Copyright 2016, ZeroC, Inc.

- `Ice::Object::ice_getHash()`
  This C++ method was deprecated in Ice 3.5.

- `IcePatch2.ChunkSize`
  `IcePatch2.Directory`
  `IcePatch2.Remove`
  `IcePatch2.Thorough`
  These properties were deprecated in previous Ice releases and replaced by properties that use the prefix `IcePatch2Client`.

- `Glacier2.AddSSLContext`
  This property was deprecated in Ice 3.3.1 and replaced by `Glacier2.AddConnectionContext`.

The following components were removed in Ice 3.6:

- Qt SQL database plug-ins for IceGrid and IceStorm, which were deprecated in Ice 3.5.
  Freeze is now the only persistence mechanism for these services.

Deprecated APIs

This section discusses APIs and components that are now deprecated. These APIs will be removed in a future Ice release, therefore we encourage you to update your applications and eliminate the use of these APIs as soon as possible.

The following APIs were deprecated in Ice 3.6:

- `Ice.ACM.Client`
  `Ice.ACM.Server`
  Use `Ice.ACM.Client.Timeout` and `Ice.ACM.Server.Timeout` instead. See Active Connection Management for more information.

- `IceSSL.CertAuthFile`
  Use the new property `IceSSL.CAs` to specify the path name of a PEM file containing the Root Certificate Authorities.

- `IceSSL.CertAuthDir` (OpenSSL only)
  Use the new property `IceSSL.CAs` to specify the path name of a directory containing the Root Certificate Authorities.

- `IceSSL.KeyFile`
  Use `IceSSL.CertFile` to configure the IceSSL identity using a PKCS12 file.

- `IceSSL.ImportCert.*` (.NET only)
  This property caused the IceSSL plug-in to import a certificate into a Windows certificate store. Going forward, users will need to install certificates in a store using Windows-provided tools if necessary. However, it is now possible to use trusted CA certificates from a file without loading them into a store. See `IceSSL.CAs` for more information.

- `IceSSL.PersistKeySet` (.NET only)
  Only used by the now-deprecated `IceSSL.ImportCert` property.

- `IceSSL.KeySet` (.NET only)
  Use `IceSSL.CertStoreLocation` instead.

- `IceSSL::Certificate::verify` (C++ only)
  The method `verify(const PublicKeyPtr&)` has been deprecated. The new method `verify(const CertificatePtr&)` ta kes its place.

- `IceBox.InstanceName`
  `IceBox.ServiceManager.`*`AdapterProperty`*
  These properties are no longer necessary because their functionality is provided by the service manager administrative facet.

- `clr:collection`
  `Ice.CollectionBase`
  `Ice.DictionaryBase`
  The `clr:collection` metadata tag, along with the C# classes `Ice.CollectionBase` and `Ice.DictionaryBase`, were originally provided for backward compatibility with Ice versions prior to 3.3 and are now deprecated. Existing applications should migrate their code to use the standard C# collection types.

- `Ice::StringConverterPlugin` (C++ only)
  The `StringConvererPlugin` base class is deprecated. You should use a regular `Ice::Plugin` to install your string

Copyright 2016, ZeroC, Inc.

converter(s).

- `Ice::CollocatedOptimizationException`
  This exception is no longer used now that collocated invocations have become much more flexible.

- `Ice.BatchAutoFlush`
  This property controlled whether the Ice run time would automatically flush batch requests for a connection after enough requests had been queued to reach the limit established by `Ice.MessageSizeMax`. We're deprecating this property in favor of a new one, `Ice.BatchAutoFlushSize`, whose value Ice now uses as the limit for automatic flushing. If your application sets `Ice.BatchAutoFlush=0` to disable automatic flushing, you can achieve the same behavior by setting `Ice.BatchAutoFlushSize=0`.

Visual C++ compiler warnings

The Ice 3.5 header files downgrade the following Visual C++ warnings to level 4:

C++

```
#   pragma warning( 4 : 4250 ) // ... : inherits ... via dominance
#   pragma warning( 4 : 4251 ) // class ... needs to have dll-interface to be used
by clients of class ...
```

This downgrade affects any file that includes these Ice header files.

Ice 3.6 no longer disables or downgrades any warning in your C++ code. As a result, when upgrading to Ice 3.6, your Ice application may produce compiler warnings that were not reported before. To eliminate these warnings, you can modify your source code, add `pragmas` or disable these warnings in your Visual Studio projects.

Copyright 2016, ZeroC, Inc.

Upgrading your Application from Ice 3.4

In addition to the information provided in Upgrading your Application from Ice 3.5, users who are upgrading from Ice 3.4 should also review this page.

On this page:

- Slicing behavior in Ice 3.5
- Slice language changes in Ice 3.5
- C++ changes in Ice 3.5
- Java changes in Ice 3.5
- Python changes in Ice 3.5
- PHP changes in Ice 3.5
- IPv6 changes in Ice 3.5
- Migrating IceGrid databases from Ice 3.4
- Migrating IceStorm databases from Ice 3.4
- Migrating Freeze databases from Ice 3.4
    - Upgrading to the Berkeley DB 5.3 format
    - Freeze maps and the 1.1 encoding
- Migrating Android applications from Ice 3.4
- Changed APIs in Ice 3.5
- Removed APIs in Ice 3.5
- Deprecated APIs in Ice 3.5

Slicing behavior in Ice 3.5

In version 1.0 of the Ice encoding, Slice classes and user exceptions are always marshaled in a format that allows a receiver to "slice" an unknown derived type to a known base type. An application that uses Ice 3.4 or earlier may depend on this behavior. For example, the slicing feature makes it possible for a sender to evolve independently of a receiver, adding more derived types that the receiver does not understand without causing the receiver to fail should it encounter one of these new types.

Version 1.1 of the Ice encoding introduces two formats for classes and exceptions. The compact format, which is used by default, does not support the slicing feature, meaning an application that relies on this feature may begin to receive `NoObjectFactoryException` or `UnknownUserException` errors after upgrading to Ice 3.5. There are several possible solutions:

- Continue using version 1.0 of the Ice encoding by setting `Ice.Default.EncodingVersion` to `1.0` in all senders and receivers.
- If you wish to take advantage of other features offered by encoding version 1.1, you can make the sliced format the default by setting `Ice.Default.SlicedFormat` to 1 in all senders and receivers.
- Annotate your interfaces or operations with metadata so that the sliced format is only used where necessary.

Slice language changes in Ice 3.5

The term `optional` is now a Slice keyword in Ice 3.5, which means existing Slice definitions that use this term as an identifier will generate an error. To continue using `optional` as an identifier, you can use the Slice syntax for escaping keywords:

Slice

```
struct Example
{
    bool \optional;
};
```

C++ changes in Ice 3.5

The `slice2cpp` translator no longer generates C++ comparison operators for all Slice structures. Comparison operators are only generated for structures that qualify as a legal dictionary key type. To generate a comparison operator for a Slice structure that doesn't qualify as such, you can prefix the `["cpp:comparable"]` metadata to your Slice struct definition.

Copyright 2016, ZeroC, Inc.

Java changes in Ice 3.5

The generated code and supporting classes for the various Ice services are now provided as separate JAR files:

JAR File	Description
`Freeze.jar`	Freeze classes
`Glacier2.jar`	Glacier2 generated code and helper classes
`Ice.jar`	The core Ice run time, including the IceSSL plugin
`IceBox.jar`	IceBox server, generated code, admin utility
`IceGrid.jar`	IceGrid generated code
`IcePatch2.jar`	IcePatch2 generated code
`IceStorm.jar`	IceStorm generated code

Also note that the manifest in `Freeze.jar` no longer contains a reference to `db.jar`, so you may need to update your class path to include `db.jar`.

The Eclipse Plug-in still includes only `Ice.jar` by default, but you can easily include additional JAR files by changing the Slice2Java properties for your project.

Python changes in Ice 3.5

Ice 3.5 adds support for Python 3. Developers who are migrating existing Ice applications from Python 2 to Python 3 should be aware of the following changes that affect the Python language mapping:

- One of the most significant changes in Python 3 is its unification of the unicode and string types into just one type, `str`. Python developers must convert any existing unicode literals and variables into their string equivalents. For Ice developers, a benefit of this unification is that all string values sent "over the wire" are now sent and received as unicode strings. With Python 2.x, the Ice run time behavior remains unchanged: string and unicode values are both accepted as inputs, but only string values are returned as outputs.
- The Python 3 mapping for the Slice type `sequence<byte>` now defaults to the new Python type `bytes`. A string value is no longer accepted as a legal value for `sequence<byte>` with Python 3.

PHP changes in Ice 3.5

For compatibility with newer PHP releases, the Ice extension no longer supports "call-time" reference arguments when calling `Ice_initialize` and `Ice_createProperties`. For example, existing applications may invoke `Ice_initialize` with an explicit reference argument as follows:

PHP

```
$communicator = Ice_initialize(&$argv);
```

The proper way to call this function using Ice 3.5 is shown below:

PHP

```
$communicator = Ice_initialize($argv);
```

Ice-specific arguments are now removed from the given argument vector.

Copyright 2016, ZeroC, Inc.

IPv6 changes in Ice 3.5

IPv6 is now enabled by default on all platforms. You can disable the use of IPv6 in Ice by setting the `Ice.IPv6` property to `0`, which may be necessary in environments that lack IPv6 support. The new property `Ice.PreferIPv6Address` determines whether IPv4 or IPv6 addresses take precedence when resolving host names.

Migrating IceGrid databases from Ice 3.4

Ice 3.5 supports the migration of IceGrid databases from Ice 3.3 and from Ice 3.4. To migrate from earlier Ice versions, you will first need to migrate the databases to the Ice 3.3 format. If you require assistance with such migration, please contact support@zeroc.com.

To migrate, first stop the IceGrid registry you wish to upgrade.

Next, copy the IceGrid database environment to a second location:

```
$ cp -r db recovered.db
```

Locate the correct version of the Berkeley DB recovery tool (usually named `db_recover`). It is essential that you use the `db_recover` executable that matches the Berkeley DB version of your existing Ice release. For Ice 3.3, use `db_recover` from Berkeley DB 4.6. For Ice 3.4, use `db_recover` from Berkeley DB 4.8. You can verify the version of your `db_recover` tool by running it with the `-V` option:

```
$ db_recover -V
```

Now run the utility on your copy of the database environment:

```
$ db_recover -h recovered.db
```

Change to the location where you will store the database environments for IceGrid 3.5:

```
$ cd <new-location>
```

Next, run the `upgradeicegrid35.py` utility located in the `config` directory of your Ice distribution (or in `/usr/share/Ice-3.5` if using an RPM installation). The first argument is the path to the old database environment. The second argument is the path to the new database environment.

In this example we'll create a new directory `db` in which to store the migrated database environment:

```
$ mkdir db
$ upgradeicegrid35.py <path-to-recovered.db> db
```

Upon completion, the `db` directory contains the migrated IceGrid databases.

By default, the migration utility assumes that the servers deployed with IceGrid also use Ice 3.5. If your servers still use an older Ice version, you need to specify the `--server-version` command-line option when running `upgradeicegrid35.py`:

```
$ upgradeicegrid.py --server-version 3.4.2 <path-to-recovered.db> db
```

The migration utility will set the `server` descriptor attribute `ice-version` to the specified version and the IceGrid registry will generate configuration files compatible with the given version.

If you are upgrading the master IceGrid registry in a replicated environment and the slaves are still running, you should first restart the master registry in read-only mode using the `--readonly` option, for example:

```
$ icegridregistry --Ice.Config=config.master --readonly
```

Next, you can connect to the master registry with `icegridadmin` or the IceGrid administrative GUI from Ice 3.5 to ensure that the database is correct. If everything looks fine, you can shutdown and restart the master registry without the `--readonly` option.

IceGrid slaves from Ice 3.3 or 3.4 won't interoperate with the IceGrid 3.5 master. You can leave them running during the upgrade of the master to not interrupt your applications. Once the master upgrade is done, you should upgrade the IceGrid slaves to Ice 3.5 using the instruction above.

Migrating IceStorm databases from Ice 3.4

No changes were made to the database schema for IceStorm in this release. However, you still need to update your databases as described below.

Copyright 2016, ZeroC, Inc.

Migrating Freeze databases from Ice 3.4

No changes were made that would affect the content of your Freeze databases. However, we upgraded the version of Berkeley DB and the new 1.1 encoding may require that you re-create your indices or set some additional configuration if you use Freeze maps. If you only use Freeze evictors, you only need to upgrade your databases to the new Berkeley DB format. See Encoding Version 1.1.

Upgrading to the Berkeley DB 5.3 format

When upgrading to Ice 3.5, you must upgrade your database to the Berkeley DB 5.3 format. The only change that affects Freeze is the format of Berkeley DB's log file.

The instructions below assume that the database environment to be upgraded resides in a directory named `db` in the current working directory. For a more detailed discussion of database migration, please refer to the Berkeley DB Upgrade Process.

To migrate your database:

1. Shut down the old version of the application.
2. Make a backup copy of the database environment:

```
> cp -r db backup.db      (Unix)
> xcopy /E db backup.db   (Windows)
```

3. Locate the correct version of the Berkeley DB recovery tool (usually named `db_recover`). It is essential that you use the `db_recover` executable that matches the Berkeley DB version of your existing Ice release. For Ice 3.4, use `db_recover` from Berkeley DB 4.8. You can verify the version of your `db_recover` tool by running it with the `-V` option:

```
> db_recover -V
```

4. Use the `db_recover` tool to run recovery on the database environment:

```
> db_recover -h db
```

5. Recompile and install the new version of the application.
6. Force a checkpoint using the `db_checkpoint` utility. Note that you must use the `db_checkpoint` utility from Berkeley DB 5.3 when performing this step.

```
> db_checkpoint -1 -h db
```

7. Restart the application.

Freeze maps and the 1.1 encoding

The majority of Freeze maps will work out of the box with the new 1.1 encoding and without additional configuration. However there are a few cases where you will be required to take special action or add some configuration:

- If your Freeze maps use keys with Slice enumerations having more than 127 elements, the encoding of those keys will be incompatible with the 1.1 encoding. You will need to set the `Freeze.DbEnv.<env-name>.EncodingVersion` property to `1.0` to ensure that you decode the keys with the 1.0 encoding and not with the default 1.1 encoding.
- If your Freeze maps use indices and those indices use Slice classes, or use Slice enumerations having more than 127 elements, you have two options:
    1. Re-create the indices with the 1.1 encoding. To recreate your indices, you have to recreate the Freeze maps using the `recreate` method generated on the Freeze map classes by `slice2freezej`. See the Freeze Maps documentation for more information.
    2. Set the `Freeze.DbEnv.<env-name>.EncodingVersion` property to `1.0` to ensure that Freeze still uses the 1.0 encoding for your indices.

Copyright 2016, ZeroC, Inc.

- If you use indices with Java Freeze maps and the map value uses classes, you have to recreate the indices because of a bug in the encoding of the Ice 3.4 indices. This only affects Java Freeze maps; C++ Freeze maps are not affected by this issue.

If you configure your Freeze database environment to use the 1.0 encoding, you will not be able to take advantage of the new features provided with the 1.1 encoding, such as optional values and the new formats for classes.

Migrating Android applications from Ice 3.4

Prior versions of the Ice plug-in for Eclipse (including the version released for Ice 3.5b) created a workspace variable named `ICE_HOME` that referred to the Ice installation directory configured in Eclipse's *Preferences* dialog. After adding the Slice2Java builder to a project, the builder updated the project's build path to include Ice JAR files relative to `ICE_HOME`. For example, `Ice.jar` was configured as the library reference `ICE_HOME/lib/Ice.jar`.

As of Ice 3.5.0, the Eclipse plug-in no longer uses the `ICE_HOME` variable. Instead, it now uses the workspace variable `ICE_JAR_HOME` to refer to the subdirectory containing `Ice.jar`, equivalent to the previous setting of `ICE_HOME/lib`. Consequently, a new project's build path now includes the library reference `ICE_JAR_HOME/Ice.jar`.

The plug-in does not attempt to modify the build path of an existing Ice for Android project to use `ICE_JAR_HOME`, nor does it remove `ICE_HOME` from the workspace. An existing project can continue to use the `ICE_HOME` variable in its build path, but moving or copying the project to a new workspace where `ICE_HOME` is no longer defined means the project's build path must be updated.

Changed APIs in Ice 3.5

This section describes APIs whose semantics have changed, potentially in ways that are incompatible with previous releases.

The following APIs were changed in Ice 3.5:

- The dynamic invocation and dispatch API now requires encapsulations for input parameters and output results. For example, when calling `ice_invoke`, the value for `inParams` must be an encapsulation of the input parameters:

  C++

  ```
  Ice::OutputStreamPtr out = Ice::createOutputStream();
  out->startEncapsulation();
  // marshal parameters
  out->endEncapsulation();
  Ice::ByteSeq inParams;
  out->finished(inParams);
  // invoke operation...
  ```

  Similarly, the byte sequence containing the results must also be an encapsulation:

  C++

  ```
  Ice::ByteSeq results;
  // invoke operation...
  Ice::InputStreamPtr in = Ice::wrapInputStream(results);
  in->startEncapsulation();
  // unmarshal results
  in->endEncapsulation();
  ```

- The streaming API uses new methods to mark the beginning and end of objects and exceptions: `startObject`/`endObject` and `startException`/`endException`.

Copyright 2016, ZeroC, Inc.

Removed APIs in Ice 3.5

This section describes APIs that were deprecated in a previous release and have now been removed. Your application may no longer compile successfully if it relies on one of these APIs.

The following APIs were removed in Ice 3.5:

- `Ice.Default.CollocationOptimization`
  Use `Ice.Default.CollocationOptimized` instead.

- *`proxy`*`.CollocationOptimization`
  Use *`proxy`*`.CollocationOptimized` instead.

- *`adapter`*`.RegisterProcess`
  This property caused the Ice run time to register a proxy with the locator registry that allowed the process to be shut down remotely. The new administrative facility has replaced this functionality.

- `Ice.ServerId`
  As with *`adapter`*`.RegisterProcess`, this property was used primarily for IceGrid integration and has been replaced by a similar mechanism in the administrative facility.

- `Ice.Trace.Location`
  This property has been replaced by `Ice.Trace.Locator`.

- `Glacier2.Admin` and `IcePatch2.Admin`
  These are the names of administrative object adapters in Glacier2 and IcePatch2, respectively. The functionality offered by these object adapters has been replaced by that of the administrative facility, therefore these adapters (and their associated configuration properties) are no longer necessary.

- `Ice.Util.generateUUID()`
  In Java use `java.util.UUID.randomUUID().toString()`. In C# use `System.Guid.NewGuid.ToString()`.

- `Object.ice_hash()`
  `ObjectPrx.ice_getHash()`
  `ObjectPrx.ice_toString()`
  Ice no longer defines hash methods for objects and proxies. To convert a proxy into a string, use the standard platform methods: `toString` in Java, `ToString` in C#.

Deprecated APIs in Ice 3.5

This section discusses APIs and components that are now deprecated. These APIs will be removed in a future Ice release, therefore we encourage you to update your applications and eliminate the use of these APIs as soon as possible.

The following APIs were deprecated in Ice 3.5:

- `Stats` facility
  This functionality is now provided by the Instrumentation facility and the Metrics facet.

The following components were deprecated in Ice 3.5:

- Qt SQL database plug-ins for IceGrid and IceStorm
  Freeze is now the only supported persistence mechanism for these services.

Copyright 2016, ZeroC, Inc.

Upgrading your Application from Ice 3.3

In addition to the information provided in Upgrading your Application from Ice 3.4, users who are upgrading from Ice 3.3 should also review this page.

On this page:

- Java language mapping changes in Ice 3.4
  - Metadata
  - Dictionaries
  - Request Contexts
  - Enumerations
- Changes to the Java API for Freeze maps in Ice 3.4
  - General changes to Freeze maps in Java
  - Enhancements to Freeze maps in Java
  - Backward compatibility for Freeze maps in Java
  - Finalizers in Freeze
- Freeze packaging changes in Ice 3.4
- PHP changes in Ice 3.4
  - Static translation in PHP
  - Deploying a PHP application
  - Using communicators in PHP
  - Using registered communicators in PHP
  - PHP configuration
  - PHP namespaces
  - Run-time exceptions in PHP
  - Downcasting in PHP
  - Other API changes for PHP
- Thread pool changes in Ice 3.4
- IceSSL changes in Ice 3.4
- Migrating IceStorm and IceGrid databases from Ice 3.3
- Migrating Freeze databases from Ice 3.3
- Removed APIs in Ice 3.4.0
- Deprecated APIs in Ice 3.4.0

Java language mapping changes in Ice 3.4

The Java2 language mapping, which was deprecated in Ice 3.3, is no longer supported. The Slice compiler and Ice API now use the Java5 language mapping exclusively, therefore upgrading to Ice 3.4 may require modifications to your application's source code. The subsections below discuss the language mapping features that are affected by this change and describe how to modify your application accordingly.

Metadata

The global metadata directives `java:java2` and `java:java5` are no longer supported and should be removed from your Slice files. The Slice compiler now emits a warning about these directives.

Support for the portable metadata syntax has also been removed. This syntax allowed Slice definitions to define custom type metadata that the Slice compiler would translate to match the desired target mapping. For example:

Slice

```
["java:type:{java.util.ArrayList}"] sequence<String> StringList;
```

The braces surrounding the custom type `java.util.ArrayList` directed the Slice compiler to use `java.util.ArrayList<String>` in the Java5 mapping and `java.util.ArrayList` in the Java2 mapping.

All uses of the portable metadata syntax must be changed to use the corresponding Java5 equivalent.

Dictionaries

Copyright 2016, ZeroC, Inc.

Now that Slice dictionary types use the Java5 mapping, recompiling your Slice files and your application may cause the Java compiler to emit "unchecked" warnings. This occurs when your code attempts to assign an untyped collection class such as `java.util.Map` to a generic type such as `java.util.Map<String, String>`. Consider the following example:

Slice

```
dictionary<string, int> ValueMap;

interface Table
{
    void setValues(ValueMap m);
};
```

A Java2 application might have used these Slice definitions as shown below:

Java

```
java.util.Map values = new java.util.HashMap();
values.put(...);

TablePrx proxy = ...;
proxy.setValues(values); // Warning
```

The call to `setValues` is an example of an unchecked conversion. We recommend that you compile your application using the compiler option shown below:

```
javac -Xlint:unchecked ...
```

This option causes the compiler to generate descriptive warnings about occurrences of unchecked conversions to help you find and correct the offending code.

Request Contexts

The Slice type for request contexts, `Ice::Context`, is defined as follows:

Slice

```
module Ice
{
    dictionary<string, string> Context;
};
```

As a dictionary, the `Context` type is subject to the same issues regarding unchecked conversions described for #Dictionaries. For example, each proxy operation maps to two overloaded methods, one that omits the trailing `Context` parameter and one that includes it:

Copyright 2016, ZeroC, Inc.

Java

```
interface TablePrx
{
    void setValues(java.util.Map<String, Integer> m);  // No context

    void setValues(java.util.Map<String, Integer> m,
                   java.util.Map<String, String> ctx);
}
```

If your proxy invocations make use of this parameter, you will need to change your code to use the generic type shown above in order to eliminate unchecked conversion warnings.

Enumerations

The Java2 language mapping for a Slice enumeration generated a class whose API differed in several ways from the standard Java5 enum type. Consider the following enumeration:

Slice

```
enum Color { red, green, blue };
```

The Java2 language mapping for `Color` is shown below:

Java

```
public final class Color
{
    // Integer constants
    public static final int _red = 0;
    public static final int _green = 1;
    public static final int _blue = 2;

    // Enumerators
    public static final Color red = ...;
    public static final Color green = ...;
    public static final Color blue = ...;

    // Helpers
    public static Color convert(int val);
    public static Color convert(String val);
    public int value();

    ...
}
```

The first step in migrating to the Java5 mapping for enumerations is to modify all `switch` statements that use an enumerator. Before Java

Copyright 2016, ZeroC, Inc.

added native support for enumerations, the `switch` statement could only use the integer value of the enumerator and therefore the Java2 mapping supplied integer constants for use in `case` statements. For example, here is a `switch` statement that uses the Java2 mapping:

Java

```
Color col = ...;
switch(col.value())
{
case Color._red:
    ...
    break;
case Color._green:
    ...
    break;
case Color._blue:
    ...
    break;
}
```

The Java5 mapping eliminates the integer constants because Java5 allows enumerators to be used in `case` statements. The resulting code becomes much easier to read and write:

Java

```
Color col = ...;
switch(col)
{
case red:
    ...
    break;
case green:
    ...
    break;
case blue:
    ...
    break;
}
```

The next step is to replace any uses of the `value` or `convert` methods with their Java5 equivalents. The base class for all Java5 enumerations (`java.lang.Enum`) supplies methods with similar functionality:

Java

```
static Color[] values()          // replaces convert(int)
static Color valueOf(String val) // replaces convert(String)
int ordinal()                    // replaces value()
```

For example, here is the Java5 code to convert an integer into its equivalent enumerator:

Copyright 2016, ZeroC, Inc.

Java

```
Color r = Color.values()[0]; // red
```

Note however that the `convert(String)` method in the Java2 mapping returned null for an invalid argument, whereas the Java5 enum method `valueOf(String)` raises `IllegalArgumentException` instead.

Refer to the manual for more details on the mapping for enumerations.

Changes to the Java API for Freeze maps in Ice 3.4

The Java API for Freeze maps has been revised to used Java5 generic types and enhanced to provide additional functionality. This section describes these changes in detail and explains how to migrate your Freeze application to the API in Ice 3.4.

General changes to Freeze maps in Java

The Freeze API is now entirely type-safe, which means compiling your application against Ice 3.4 is likely to generate unchecked conversion warnings. The generated class for a Freeze map now implements the `java.util.SortedMap<K, V>` interface, where *K* is the key type and *V* is the value type. As a result, applications that relied on the untyped `SortedMap` API (where all keys and values were treated as instances of `java.lang.Object`) will encounter compiler warnings in Ice 3.4.

For example, an application might have iterated over the entries in a map as follows:

Java

```
// Old API
Object key = new Integer(5);
Object value = new Address(...);
myMap.put(key, value);
java.util.Iterator i = myMap.entrySet().iterator();
while (i.hasNext())
{
    java.util.Map.Entry e = (java.util.Map.Entry)i.next();
    Integer myKey = (Integer)e.getKey();
    Address myValue = (Address)e.getValue();
    ...
}
```

This code will continue to work, but the new API is both type-safe and self-documenting:

Copyright 2016, ZeroC, Inc.

Java

```
// New API
int key = 5;
Address value = new Address(...);
myMap.put(key, value); // The key is autoboxed to Integer.
for (java.util.Map.Entry<Integer, Address> e : myMap.entrySet())
{
    Integer myKey = e.getKey();
    Address myValue = e.getValue();
    ...
}
```

Although migrating to the new API may require some effort, the benefits are worthwhile because your code will be easier to read and less prone to defects. You can also take advantage of the "autoboxing" features in Java5 that automatically convert values of primitive types (such as `int`) into their object equivalents (such as `Integer`).

Please refer to the manual for complete details on the new API.

Enhancements to Freeze maps in Java

Java6 introduced the `java.util.NavigableMap` interface, which extends `java.util.SortedMap` to add some useful new methods. Although the Freeze map API cannot implement `java.util.NavigableMap` directly because Freeze must remain compatible with Java5, we have added the `Freeze.NavigableMap` interface to provide much of the same functionality. A generated Freeze map class implements `NavigableMap`, as do the sub map views returned by map methods such as `headMap`. The `NavigableMap` interface is described in the manual, and you can also refer to the Java6 API documentation.

Backward compatibility for Freeze maps in Java

The Freeze Map API related to indices underwent some significant changes in order to improve type safety and avoid unchecked conversion warnings. These changes may cause compilation failures in a Freeze application.

In the previous API, index comparator objects were supplied to the Freeze map constructor in a map (in Java5 syntax, this comparators map would have the type `java.util.Map<String, java.util.Comparator>`) in which the index name was the key. As part of our efforts to improve type safety, we also wanted to use the fully-specified type for each index comparator (such as `java.util.Comparator<Integer>`). However, given that each index could potentially use a different key type, it is not possible to retain the previous API while remaining type-safe.

Consequently, the index comparators are now supplied as data members of a static nested class of the Freeze map named `IndexComparators`. If your application supplied custom comparators for indices, you will need to revise your code to use `IndexComparators` instead. For example:

Copyright 2016, ZeroC, Inc.

Java

```
// Old API
java.util.Map indexComparators = new java.util.HashMap();
indexComparators.put("index", new MyComparator());
MyMap map = new MyMap(..., indexComparators);

// New API
MyMap.IndexComparators indexComparators = new MyMap.IndexComparators();
indexComparators.valueComparator = new MyComparator();
MyMap map = new MyMap(..., indexComparators);
```

We also encourage you to modify the definition of your comparator classes to use the Java5 syntax, as shown in the example below:

Java

```
// Old comparator
class IntComparator implements java.util.Comparator
{
    public int compare(Object o1, Object o2)
    {
        return ((Integer)o1).compareTo(o2);
    }
}

// New comparator
class IntComparator implements java.util.Comparator<Integer>
{
    public int compare(Integer i1, Integer i2)
    {
        return i1.compareTo(i2);
    }
}
```

The second API change that might cause compilation failures is the removal of the following methods:

Java

```
java.util.SortedMap headMapForIndex(String name, Object key);
java.util.SortedMap tailMapForIndex(String name, Object key);
java.util.SortedMap subMapForIndex(String name, Object from, Object to);
java.util.SortedMap mapForIndex(String name);
```

Again, this API cannot be retained in a type-safe fashion, therefore `slice2freezej` now generates equivalent (and type-safe) methods for each index in the Freeze map class.

Please refer to the manual for complete details on the new API.

Copyright 2016, ZeroC, Inc.

Finalizers in Freeze

In previous releases, Freeze for Java used finalizers to close objects such as maps and connections that the application neglected to close. Most of these finalizers have been removed in Ice 3.4, and the only remaining finalizers simply log warning messages to alert you to the fact that connections and iterators are not being closed explicitly. Note that, given the uncertain nature of Java finalizers, it is quite likely that the remaining finalizers will not be executed.

Freeze packaging changes in Ice 3.4

All Freeze-related classes are now stored in a separate JAR file named `Freeze.jar`. As a result, you may need to update your build scripts, deployment configuration, and run-time environment to include this additional JAR file.

PHP changes in Ice 3.4

The Ice extension for PHP has undergone many changes in this release. The subsections below describe these changes in detail. Refer to the PHP Mapping for more information about the language mapping.

Static translation in PHP

In prior releases, Slice files were deployed with the application and loaded at Web server startup by the Ice extension. Before each page request, the extension directed the PHP interpreter to parse the code that was generated from the Slice definitions.

In this release, Slice files must be translated using the new compiler `slice2php`. This change offers several advantages:

- Applications may have more opportunities to improve performance through the use of opcode caching.
- It is no longer necessary to restart the Web server when you make changes to your Slice definitions, which is especially useful during development.
- Errors in your Slice files can now be discovered in your development environment, rather than waiting until the Web server reports a failure and then reviewing the server log to determine the problem.
- The development process becomes simpler because you can easily examine the generated code if you have questions about the API or language mapping rules.
- PHP scripts can now use all of the Ice local exceptions. In prior releases, only a subset of the local exception types were available, and all others were mapped to `Ice_UnknownLocalException`. See the section Run-time exceptions in PHP below for more information.

All of the Slice files for Ice and Ice services are translated during an Ice build and available for inclusion in your application. At a minimum, you must include the file `Ice.php`:

PHP

```
require 'Ice.php';
```

`Ice.php` contains definitions for core Ice types and includes a minimal set of generated files. To use an Ice service such as IceStorm, include the appropriate generated file:

PHP

```
require 'Ice.php';
require 'IceStorm/IceStorm.php';
```


Copyright 2016, ZeroC, Inc.

Deploying a PHP application

With the transition to static code generation, you no longer need to deploy Slice files with your application. Instead, you will need to deploy the PHP code generated from your Slice definitions, along with `Ice.php`, the generated code for the Ice core, and the generated code for any Ice services your application might use.

Using communicators in PHP

In prior releases, each PHP page request could access a single Ice communicator via the `$ICE` global variable. The configuration of this communicator was derived from the profile that the script loaded via the `Ice_loadProfile` function. The communicator was created on demand when `$ICE` was first used and destroyed automatically at the end of the page request.

In this release, a PHP script must create its own communicator using an API that is similar to other Ice language mappings:

PHP

```
function Ice_initialize()
function Ice_initialize($args)
function Ice_initialize($initData)
function Ice_initialize($args, $initData)
```

`Ice_initialize` creates a new communicator using the configuration provided in the optional arguments. `$args` is an array of strings representing command-line options, and `$initData` is an instance of `Ice_InitializationData`.

An application that requires no configuration can initialize a communicator as follows:

PHP

```
$communicator = Ice_initialize();
```

More elaborate configuration scenarios are described in the section #PHP configuration below.

A script may optionally destroy its communicator:

PHP

```
$communicator->destroy();
```

At the completion of a page request, Ice by default automatically destroys any communicator that was not explicitly destroyed.

Using registered communicators in PHP

PHP applications may benefit from the ability to use a communicator instance in multiple page requests. Reusing a communicator allows the application to minimize the overhead associated with the communicator lifecycle, including such activities as opening and closing connections to Ice servers.

This release includes new APIs for registering a communicator in order to prevent Ice from destroying it automatically at the completion of a page request. For example, a session-based application can create a communicator, establish a Glacier2 session, and register the communicator. In subsequent page requests, the PHP session can retrieve its communicator instance and continue using the Glacier2 session.

The manual provides more information on this feature, and a new sample program can be found in `Glacier2/hello`.

Copyright 2016, ZeroC, Inc.

PHP configuration

Prior releases supported four INI settings in PHP's configuration file:

- `ice.config`
- `ice.options`
- `ice.profiles`
- `ice.slice`

The `ice.slice` directive is no longer supported since Slice definitions are now compiled statically. The remaining options are still supported but their semantics are slightly different. They no longer represent the configuration of a communicator; instead, they define property sets that a script can retrieve and use to initialize a communicator.

The global INI directives `ice.config` and `ice.options` configure the default property set. The `ice.profiles` directive can optionally nominate a separate file that defines any number of named profiles, each of which configures a property set.

As before, the profiles use an INI file syntax:

```
[Name1]
config=file1
options="--Ice.Trace.Network=2 ..."

[Name2]
config=file2
options="--Ice.Trace.Locator=1 ..."
```

A new directive, `ice.hide_profiles`, overwrites the value of the `ice.profiles` directive as a security measure. This directive has a default value of 1, meaning it is enabled by default.

A script can obtain a property set using the new function `Ice_getProperties`. Called without an argument (or with an empty string), the function returns the default property set:

PHP

```
$props = Ice_getProperties();
```

Alternatively, you can pass the name of the desired profile:

PHP

```
$props = Ice_getProperties("Name1");
```

The returned object is an instance of `Ice_Properties`, which supports the standard Ice API.

For users migrating from an earlier release, you can replace a call to `Ice_loadProfile` as follows:

Copyright 2016, ZeroC, Inc.

PHP

```
// PHP - Old API
Ice_loadProfile('Name1');

// PHP - New API
$initData = new Ice_InitializationData;
$initData->properties = Ice_getProperties('Name1');
$ICE = Ice_initialize($initData);
```

(Note that it is not necessary to use the symbol `$ICE` for your communicator. However, using this symbol may ease your migration to this release.)

`Ice_loadProfile` also installed the PHP definitions corresponding to your Slice types. In this release you will need to add `require` statements to include your generated code.

Finally, if you wish to manually configure a communicator, you can create a property set using `Ice_createProperties`:

PHP

```
function Ice_createProperties($args=null, $defaultProperties=null)
```

`$args` is an array of strings representing command-line options, and `$defaultProperties` is an instance of `Ice_Properties` that supplies default values for properties.

As an example, an application can configure a communicator as shown below:

PHP

```
$initData = new Ice_InitializationData;
$initData->properties = Ice_createProperties();
$initData->properties->setProperty("Ice.Trace.Network", "1");
...
$ICE = Ice_initialize($initData);
```

PHP namespaces

This release includes optional support for PHP namespaces, which was introduced in PHP 5.3. Support for PHP namespaces is disabled by default; to enable it, you must build the Ice extension from source code with `USE_NAMESPACES=yes` (see `Make.rules` or `Make.rules.mak` in the `php/config` subdirectory). Note that the extension only supports one mapping style at a time; installing a namespace-enabled version of the extension requires all Ice applications on the target Web server to use namespaces.

With namespace support enabled, you must modify your script to include a different version of the core Ice types:

Copyright 2016, ZeroC, Inc.

PHP

```
require 'Ice_ns.php'; // Namespace version of Ice.php
```

You must also recompile your Slice files using the `-n` option to generate namespace-compatible code:

```
% slice2php -n MySliceFile.ice
```

This mapping translates Slice modules into PHP namespaces instead of using the "flattened" (underscore) naming scheme. For example, `Ice_Properties` becomes `\Ice\Properties` in the namespace mapping. However, applications can still refer to global Ice functions by their traditional names (such as `Ice_initialize`) or by their namespace equivalents (`\Ice\initialize`).

Run-time exceptions in PHP

As mentioned earlier, prior releases of Ice for PHP only supported a limited subset of the standard run-time exceptions. An occurrence of an unsupported local exception was mapped to `Ice_UnknownLocalException`.

This release adds support for all local exceptions, which allows an application to more easily react to certain types of errors:

PHP

```
try
{
    $proxy->sayHello();
}
catch(Ice_ConnectionLostException $ex)
{
    // Handle connection loss
}
catch(Ice_LocalException $ex)
{
    // Handle other errors
}
```

This change represents a potential backward compatibility issue: applications that previously caught `Ice_UnknownLocalException` may need to be modified to catch the intended exception instead.

Downcasting in PHP

In prior releases, to downcast a proxy you had to invoke the `ice_checkedCast` or `ice_uncheckedCast` method on a proxy and supply a type ID:

Copyright 2016, ZeroC, Inc.

PHP

```
$hello = $proxy->ice_checkedCast("::Demo::Hello");
```

This API is susceptible to run-time errors because no validation is performed on the type ID string. For example, renaming the `Hello` interface to `Greeting` requires that you not only change all occurrences of `Demo_Hello` to `Demo_Greeting`, but also fix any type ID strings that your code might have embedded. The PHP interpreter does not provide any assistance if you forget to make this change, and you will only discover it when that particular line of code is executed and fails.

To improve this situation, a minimal class is now generated for each proxy type. The purpose of this class is to supply `checkedCast` and `uncheckedCast` static methods:

PHP

```
class Demo_HelloPrx
{
    public static function checkedCast($proxy, $facetOrCtx=null, $ctx=null);

    public static function uncheckedCast($proxy, $facet=null);
}
```

Now your application can downcast a proxy as follows:

PHP

```
$hello = Demo_HelloPrx::checkedCast($proxy);
```

You can continue to use `ice_checkedCast` and `ice_uncheckedCast` but we recommend migrating your application to the new methods.

Other API changes for PHP

This section describes additional changes to the Ice API in this release:

- The global variable `$ICE` is no longer defined. An application must now initialize its own communicator as described above.
- Removed the following communicator methods:

  PHP

  ```
  $ICE->setProperty()
  $ICE->getProperty()
  ```

  The equivalent methods are:

Copyright 2016, ZeroC, Inc.

PHP

```
$communicator->getProperties()->setProperty()
$communicator->getProperties()->getProperty()
```

- Removed the following global functions:

  PHP

  ```
  Ice_stringToIdentity()
  Ice_identityToString()
  ```

  The equivalent methods are:

  PHP

  ```
  $communicator->stringToIdentity()
  $communicator->identityToString()
  ```

- These functions have also been removed:

  PHP

  ```
  Ice_loadProfile()
  Ice_loadProfileWithArgs()
  Ice_dumpProfile()
  ```

  Refer to PHP configuration for more information.

Thread pool changes in Ice 3.4

A thread pool supports the ability to automatically grow and shrink as the demand for threads changes, within the limits set by the thread pool's configuration. In prior releases, the rate at which a thread pool shrinks was not configurable, but Ice 3.4.0 introduces the `ThreadIdleTime` property to allow you to specify how long a thread pool thread must remain idle before it terminates to conserve resources.

IceSSL changes in Ice 3.4

With the addition of the `ConnectionInfo` classes in this release, the `IceSSL::ConnectionInfo` structure has changed from a native type to a Slice class. This change has several implications for existing applications:

- As a Slice class, `IceSSL::ConnectionInfo` cannot provide the X509 certificate chain in its native form, therefore the chain is provided as a sequence of strings representing the encoded form of each certificate. You can use language-specific facilities to convert these strings back to certificate objects.

Copyright 2016, ZeroC, Inc.

- For your convenience, we have added a native subclass of `IceSSL::ConnectionInfo` called `IceSSL::NativeConnectionInfo`. This class provides the certificate chain as certificate objects.
- The `CertificateVerifier` interface now uses `NativeConnectionInfo` instead of `ConnectionInfo`. If your application configures a custom certificate verifier, you will need to modify your implementation accordingly.
- In C++, also note that `NativeConnectionInfo`is now managed by a smart pointer, therefore the signature of the certificate verifier method becomes the following:

C++

```
virtual bool verify(const IceSSL::NativeConnectionInfoPtr&) = 0;
```

- The `getConnectionInfo` helper function has been removed because its functionality has been replaced by the `Connection::getInfo`operation. For example, in prior releases a C++ application would do the following:

C++

```
Ice::ConnectionPtr con = ...
IceSSL::ConnectionInfo info = IceSSL::getConnectionInfo(con);
```

Now the application should do this:

C++

```
Ice::ConnectionPtr con = ...
IceSSL::ConnectionInfoPtr info =
IceSSL::ConnectionInfoPtr::dynamicCast(con->getInfo());
```

Alternatively, the application can downcast to the native class:

C++

```
Ice::ConnectionPtr con = ...
IceSSL::NativeConnectionInfoPtr info =
    IceSSL::NativeConnectionInfoPtr::dynamicCast(con->getInfo());
```

Migrating IceStorm and IceGrid databases from Ice 3.3

No changes were made to the database schema for IceStorm or IceGrid in this release. However, you still need to update your databases as described below.

Migrating Freeze databases from Ice 3.3

Copyright 2016, ZeroC, Inc.

No changes were made that would affect the content of your Freeze databases. However, we upgraded the version of Berkeley DB, therefore when upgrading to Ice 3.4 you must also upgrade your database to the Berkeley DB 4.8 format. The only change that affects Freeze is the format of Berkeley DB's log file.

The instructions below assume that the database environment to be upgraded resides in a directory named db in the current working directory. For a more detailed discussion of database migration, please refer to the Berkeley DB Upgrade Process.

To migrate your database:

1. Shut down the old version of the application.
2. Make a backup copy of the database environment:

```
> cp -r db backup.db      (Unix)
> xcopy /E db backup.db   (Windows)
```

3. Locate the correct version of the Berkeley DB recovery tool (usually named `db_recover`). It is essential that you use the `db_recover` executable that matches the Berkeley DB version of your existing Ice release. For Ice 3.3, use `db_recover` from Berkeley DB 4.6. You can verify the version of your `db_recover` tool by running it with the `-V`option:

```
> db_recover -V
```

4. Use the db_recover tool to run recovery on the database environment:

```
> db_recover -h db
```

5. Recompile and install the new version of the application.
6. Force a checkpoint using the `db_checkpoint` utility. Note that you must use the `db_checkpoint`utility from Berkeley DB 4.8 when performing this step.

```
> db_checkpoint -1 -h db
```

7. Restart the application.

Removed APIs in Ice 3.4.0

This section describes APIs that were deprecated in a previous release and have now been removed. Your application may no longer compile successfully if it relies on one of these APIs.

The following APIs were removed in Ice 3.4.0:

- `Glacier2.AddUserToAllowCategories`
  Use `Glacier2.Filter.Category.AcceptUser` instead.
- `Glacier2.AllowCategories`
  Use `Glacier2.Filter.Category.Accept` instead.
- `Ice.UseEventLog`
  Ice services (applications that use the C++ class `Ice::Service`) always use the Windows event log by default.
- `Communicator::setDefaultContext`
- `Communicator::getDefaultContext`
- `ObjectPrx:ice_defaultContext`

Copyright 2016, ZeroC, Inc.

Use the communicator's implicit request context instead.

- `nonmutating` keyword
  This keyword is no longer supported.
- `Freeze.UseNonmutating`
  Support for this property was removed along with the `nonmutating` keyword.
- `Ice::NegativeSizeException`
  The run time now throws `UnmarshalOutOfBoundsException` or `MarshalException` instead.
- `slice2docbook`
  This utility is no longer included in Ice.
- `Ice::AMD_Array_Object_ice_invoke`
  A new overloading of `ice_response` in the `AMD_Object_ice_invoke` class makes `AMD_Array_Object_ice_invoke` obsolete.
- Java2 mapping
  The Java2 mapping is no longer supported. Refer to Java language mapping changes in Ice 3.4 for more information.

Deprecated APIs in Ice 3.4.0

This section discusses APIs and components that are now deprecated. These APIs will be removed in a future Ice release, therefore we encourage you to update your applications and eliminate the use of these APIs as soon as possible.

The following APIs were deprecated in Ice 3.4.0:

- Asynchronous Method Invocation (AMI) interface
  The AMI interface in Ice 3.3 and earlier is now deprecated for C++, Java, and C#.
- `Glacier2.AddSSLContext`
  Replaced by `Glacier2.AddConnectionContext`.
- Standard platform methods should be used instead of the following:

  Java

  ```
  Ice.Object.ice_hash()         // Use hashCode
  Ice.ObjectPrx.ice_getHash()   // Use hashCode
  Ice.ObjectPrx.ice_toString()  // Use toString
  ```

  In Java, use `hashCode` and `toString`. In C#, use `GetHashCode` and `ToString`. In Ruby, use `hash` instead of `ice_getHash`.
- `Ice.Util.generateUUID()`
  In Java use `java.util.UUID.randomUUID().toString()`. In C# use `System.Guid.NewGuid.ToString()`.

Copyright 2016, ZeroC, Inc.

Upgrading your Application from Ice 3.2 or Earlier Releases

In addition to the information provided in Upgrading your Application from Ice 3.3, users who are upgrading from Ice 3.2 or earlier should also review this page.

On this page:

- Migrating IceStorm databases from Ice 3.2
- Migrating IceGrid databases from Ice 3.2
- Migrating Freeze databases from Ice 3.2
- Removed APIs in Ice 3.3
  - Thread per connection
  - .NET metadata
  - C++
  - Java
  - .NET
  - Python
  - General
  - Ice.LoggerPlugin
- Deprecated APIs in Ice 3.3
  - Sequences as dictionary keys
  - LocalObject
  - Ice.Trace.Location
  - Ice.Default.CollocationOptimization
  - <Adapter>.RegisterProcess
  - Ice.ServerId
  - Glacier2.Admin and IcePatch2.Admin

Migrating IceStorm databases from Ice 3.2

Ice 3.4 supports the migration of IceStorm databases from Ice 3.1 and from Ice 3.2. Migration from earlier Ice versions may work, but is not officially supported. If you require assistance with such migration, please contact support@zeroc.com.

To migrate, first stop your IceStorm servers.

Next, copy the IceStorm database environment to a second location:

```
$ cp -r db recovered.db
```

Locate the correct version of the Berkeley DB recovery tool (usually named `db_recover`). It is essential that you use the `db_recover` executable that matches the Berkeley DB version of your existing Ice release. For Ice 3.1, use `db_recover` from Berkeley DB 4.3.29. For Ice 3.2, use `db_recover` from Berkeley DB 4.5. You can verify the version of your `db_recover` tool by running it with the `-V` option:

```
$ db_recover -V
```

Now run the utility on your copy of the database environment:

```
$ db_recover -h recovered.db
```

Change to the location where you will store the database environments for IceStorm 3.4:

Copyright 2016, ZeroC, Inc.

```
$ cd <new-location>
```

Next, run the `icestormmigrate` utility. The first argument is the path to the old database environment. The second argument is the path to the new database environment.

In this example we'll create a new directory `db` in which to store the migrated database environment:

```
$ mkdir db
$ icestormmigrate <path-to-recovered.db> db
```

Upon completion, the `db` directory contains the migrated IceStorm databases.

Migrating IceGrid databases from Ice 3.2

Ice 3.4 supports the migration of IceGrid databases from Ice 3.1 and from Ice 3.2. Migration from earlier Ice versions may work, but is not officially supported. If you require assistance with such migration, please contact support@zeroc.com.

To migrate, first stop the IceGrid registry you wish to upgrade.

Next, copy the IceGrid database environment to a second location:

```
$ cp -r db recovered.db
```

Locate the correct version of the Berkeley DB recovery tool (usually named `db_recover`). It is essential that you use the `db_recover` executable that matches the Berkeley DB version of your existing Ice release. For Ice 3.1, use `db_recover` from Berkeley DB 4.3.29. For Ice 3.2, use `db_recover` from Berkeley DB 4.5. You can verify the version of your `db_recover` tool by running it with the `-V` option:

```
$ db_recover -V
```

Now run the utility on your copy of the database environment:

```
$ db_recover -h recovered.db
```

Change to the location where you will store the database environments for IceGrid 3.4:

```
$ cd <new-location>
```

Next, run the `upgradeicegrid.py` utility located in the `config` directory of your Ice distribution (or in `/usr/share/Ice-3.4.1` if using an RPM installation). The first argument is the path to the old database environment. The second argument is the path to the new database environment.

In this example we'll create a new directory `db` in which to store the migrated database environment:

Copyright 2016, ZeroC, Inc.

```
$ mkdir db
$ upgradeicegrid.py <path-to-recovered.db> db
```

Upon completion, the `db` directory contains the migrated IceGrid databases.

By default, the migration utility assumes that the servers deployed with IceGrid also use Ice 3.4. If your servers still use an older Ice version, you need to specify the `--server-version` command-line option when running `upgradeicegrid.py`:

```
$ upgradeicegrid.py --server-version 3.2.1 <path-to-recovered.db> db
```

The migration utility will set the `server` descriptor attribute `ice-version` to the specified version and the IceGrid registry will generate configuration files compatible with the given version.

If you are upgrading the master IceGrid registry in a replicated environment and the slaves are still running, you should first restart the master registry in read-only mode using the `--readonly` option, for example:

```
$ icegridregistry --Ice.Config=config.master --readonly
```

Next, you can connect to the master registry with `icegridadmin` or the IceGrid administrative GUI to ensure that the database is correct. If everything looks fine, you can shutdown and restart the master registry without the `--readonly` option.

Migrating Freeze databases from Ice 3.2

No changes were made that would affect the content of your Freeze databases. However, we upgraded the version of Berkeley DB, therefore when upgrading to Ice 3.4 you must also upgrade your database to the Berkeley DB 4.8 format. The only change that affects Freeze is the format of Berkeley DB's log file.

The instructions below assume that the database environment to be upgraded resides in a directory named `db` in the current working directory. For a more detailed discussion of database migration, please refer to the Berkeley DB Upgrade Process.

To migrate your database:

1. Shut down the old version of the application.
2. Make a backup copy of the database environment:

   ```
   > cp -r db backup.db      (Unix)
   > xcopy /E db backup.db   (Windows)
   ```

3. Locate the correct version of the Berkeley DB recovery tool (usually named `db_recover`). It is essential that you use the `db_recover` executable that matches the Berkeley DB version of your existing Ice release. For Ice 3.1, use `db_recover` from Berkeley DB 4.3.29. For Ice 3.2, use `db_recover` from Berkeley DB 4.5. You can verify the version of your `db_recover` tool by running it with the `-V`option:

   ```
   > db_recover -V
   ```

4. Use the `db_recover`tool to run recovery on the database environment:

Copyright 2016, ZeroC, Inc.

```
> db_recover -h db
```

5. Recompile and install the new version of the application.
6. Force a checkpoint using the `db_checkpoint` utility. Note that you must use the `db_checkpoint`utility from Berkeley DB 4.8 when performing this step.

```
> db_checkpoint -1 -h db
```

7. Restart the application.

Removed APIs in Ice 3.3

This section describes APIs that were deprecated in a previous release and have been removed in Ice 3.3. Your application may no longer compile successfully if it relies on one of these APIs.

Please refer to Removed APIs in Ice 3.4.0 for information on APIs that were removed in Ice 3.4.

Thread per connection

The primary purpose of this concurrency model was to serialize the requests received over a connection, either because the application needed to ensure that requests are dispatched in the order they are received, or because the application did not want to implement the synchronization that might be required when using the thread pool concurrency model.

Another reason for using the thread-per-connection concurrency model is that it was required by the IceSSL plug-ins for Java and C#. This requirement has been eliminated.

The ability to serialize requests is now provided by the thread pool and enabled via a new configuration property:

`<threadpool>.Serialize=1`

Please refer to the manual for more details on this feature.

Aside from the potential semantic changes involved in migrating your application to the thread pool concurrency model, other artifacts of thread-per-connection may be present in your application and must be removed:

- The configuration properties `Ice.ThreadPerConnection` and `<proxy>.ThreadPerConnection`
- The proxy methods `ice_threadPerConnection` and `ice_isThreadPerConnection`

.NET metadata

The metadata directive `cs:collection` is no longer valid. Use `["clr:collection"]` instead.

C++

The following C++ methods have been removed:

- `Application::main(int, char*[], const char*, const Ice::LoggerPtr&)`

  Use `Application::main(int, char*[], const InitializationData&)` instead.

- `initializeWithLogger`
- `initializeWithProperties`
- `initializeWithPropertiesAndLogger`

  Use `initialize(int, char*[], const InitializationData&)` instead.

Copyright 2016, ZeroC, Inc.

- `stringToIdentity`
- `identityToString`

  Use the equivalent operations on `Communicator`.

Java

The following methods have been removed:

- `Application.main(String, String[], String, Logger)`

  Use `Application.main(String, String[], InitializationData)` instead.

- `initializeWithLogger`
- `initializeWithProperties`
- `initializeWithPropertiesAndLogger`

  Use `initialize(String[], InitializationData)` instead.

.NET

The following methods have been removed:

- `Application.main(string, string[], string, Logger)`

  Use `Application.main(string, string[], InitializationData)` instead.

- `initializeWithLogger`
- `initializeWithProperties`
- `initializeWithPropertiesAndLogger`

  Use `initialize(ref string[], InitializationData)` instead.

Python

The following methods have been removed:

- `initializeWithLogger`
- `initializeWithProperties`
- `initializeWithPropertiesAndLogger`

  Use `initialize(args, initializationData)` instead.

- `stringToIdentity`
- `identityToString`
  Use the equivalent operations on `Communicator`.

General

The following methods have been removed:

- `ice_hash`
- `ice_communicator`
- `ice_collocationOptimization`
- `ice_connection`
  These proxy methods were replaced by ones of the form `ice_get...`, such as `ice_getCommunicator`. `ice_collocationOptimization` is now `ice_getCollocationOptimized`.

- `ice_newIdentity`
- `ice_newContext`
- `ice_newFacet`
- `ice_newAdapterId`
- `ice_newEndpoints`
  These proxy methods were replaced by ones that do not use `new` in their names. For example, `ice_newIdentity` was replaced

Copyright 2016, ZeroC, Inc.

by `ice_identity`.

`Ice.LoggerPlugin`

This property provided a way to install a custom logger implementation. It has been replaced by a more generalized facility for installing custom loggers.

Deprecated APIs in Ice 3.3

This section discusses APIs and components that are deprecated in Ice 3.3. These APIs will be removed in a future Ice release, therefore we encourage you to update your applications and eliminate the use of these APIs as soon as possible.

Please refer to Deprecated APIs in Ice 3.4.0 for information on APIs that were deprecated in Ice 3.4.

Sequences as dictionary keys

The use of sequences, and structures containing sequences, as the key type of a Slice dictionary is now deprecated.

`LocalObject`

The mappings for the `LocalObject` type have changed in Java, .NET, and Python. The new mappings are shown below:

Java	`java.lang.Object`
.NET	`System.Object`
Python	`object`

The types `Ice.LocalObject` and `Ice.LocalObjectImpl` are deprecated.

`Ice.Trace.Location`

This property has been replaced by `Ice.Trace.Locator`.

`Ice.Default.CollocationOptimization`

This property, as well as the corresponding proxy property, have been replaced by `Ice.Default.CollocationOptimized` and `<proxy>.CollocationOptimized`, respectively.

`<Adapter>.RegisterProcess`

This property caused the Ice run time to register a proxy with the locator registry (e.g., IceGrid) that allowed the process to be shut down remotely. The new administrative facility has replaced this functionality.

`Ice.ServerId`

As with `<Adapter>.RegisterProcess`, this property was used primarily for IceGrid integration and has been replaced by a similar mechanism in the administrative facility.

`Glacier2.Admin` and `IcePatch2.Admin`

These are the names of administrative object adapters in Glacier2 and IcePatch2, respectively. The functionality offered by these object adapters has been replaced by that of the administrative facility, therefore these adapters (and their associated configuration properties) are deprecated.

Copyright 2016, ZeroC, Inc.

Platform Notes for Ice 3.6

This page provides platform-specific notes for Ice 3.6.

On this page:

- Android notes
  - Socket connection issue in Android emulator
- Java notes
  - IceGrid Admin GUI
  - SSL issues in Java
    - Anonymous Diffie Hellman ciphersuites
    - Entropy pool causes hangs
  - IPv6 hang in Java
- SLES 12 notes
  - OpenJDK 1.7.0
- Windows notes
  - Anti-Virus when testing WebSocket
  - Multicast and Windows Store Apps

Android notes

Socket connection issue in Android emulator

An Ice application that attempts to connect to an invalid address can generate "spurious wakeup" messages in logcat when running under an emulator for Android 4.2 or later. This issue does not occur on hardware devices, however you can still experience lengthy delays before receiving an exception. As a defensive measure, it is a good idea to always set reasonable timeouts on your proxies to avoid unexpected delays.

Java notes

IceGrid Admin GUI

The IceGrid graphical administrative utility can be used with Java7 or Java8, however the metrics graph feature is only available when running with JavaFX.

A JavaFX bug can generate `NullPointerException` errors when using drag & drop in the GUI. This bug is fixed in Java7u10.

SSL issues in Java

Anonymous Diffie Hellman ciphersuites

Recent versions of Java7 and Java8 require low-strength ADH ciphersuites to be disabled when using TLS 1.0. It is no longer sufficient to use this IceSSL configuration:

```
IceSSL.Ciphers=NONE (DH_anon)
IceSSL.VerifyPeer=0
```

We recommend using this setting instead:

```
IceSSL.Ciphers=NONE (DH_anon.*AES)
IceSSL.VerifyPeer=0
```

Entropy pool causes hangs

When using the Ice for Java SSL plug-in (IceSSL), you may experience occasional hangs. The most likely reason is that your system's

Copyright 2016, ZeroC, Inc.

entropy pool is empty. If you have sufficient system privileges, you can solve this issue by editing the file *java.home*`/jre/lib/security/java.security` and changing it to use `/dev/urandom` instead of `/dev/random`. If you do not have permission to modify the security file, you can also use the command-line option shown below:

```
java -Djava.security.egd=file:/dev/urandom MyClass ...
```

IPv6 hang in Java

On systems with IPv6 enabled, you may experience occasional hangs the first time an Ice object adapter is activated within a JVM. A work-around is to disable IPv6 support by setting the Java property `java.net.preferIPv4Stack` to true. For example:

```
java -Djava.net.preferIPv4Stack=true MyClass ...
```

SLES 12 notes

OpenJDK 1.7.0

We do not recommend using OpenJDK 1.7.0_85 on SLES 12. The `IceBox/configuration` test fails with an unexpected `java.io.IOException` exception from the selector implementation when run with SSL. The Oracle and IBM JDKs do not exhibit this issue and should be used for Java applications on SLES.

Windows notes

Windows users should be aware of the following issues.

Anti-Virus when testing WebSocket

If your anti-virus intercepts `http` traffic on localhost, a number of tests in the Ice test suite may fail due to timeouts when running the test suite with the `ws` (WebSocket) protocol. We reproduced this problem with a popular anti-virus program, NOD32 from ESET. The work-around is to disable the "web access protection" feature of the anti-virus entirely or at least for localhost (127.0.0.1).

Multicast and Windows Store Apps

Network isolation for Windows Store Apps blocks multicast datagrams on the loopback interface. As a result, IceDiscovery and IceLocatorDiscovery won't be able to discover peers listening on the loopback interface.

Copyright 2016, ZeroC, Inc.

Using the Windows Binary Distribution

This page provides important information for users of the Ice binary distribution on Windows platforms.

On this page:

- Overview of the Windows binary distribution
- Setting up your Windows environment to use Ice
  - C++
  - .NET
  - Java
  - PHP
- Using the sample programs on Windows
- Configuration files for IceGrid and Glacier2 services
- Starting IceGrid Admin on Windows
- Unattended installation
- Ice installer registry key
- Third-party packages for Windows

Overview of the Windows binary distribution

The Ice binary distribution for Windows provides the Ice run-time and SDKs for the following programming languages and platforms:

- for C++, using Visual Studio 2012, Visual Studio 2013 or Visual Studio 2015
- for .NET, using Visual Studio 2012, Visual Studio 2013 or Visual Studio 2015
- for Java, using Java 7 or Java 8
- for Android, using Android 4.2 or later
- for PHP, using PHP 5.6

You only need the development environment for your target programming language to use this distribution. For example, if you want to build Ice applications in Java, you need to install a JDK, but do not need to install Visual Studio or PHP.

Ice Builders
You should also install the Ice Builder for your development environment, such as the Ice Builder for Visual Studio or the Ice Builder for Gradle. These Ice Builders are not included in the Windows binary distribution, and typically support several versions of Ice.

When upgrading to a new version of Ice, please remember to update the Ice Home configuration of the Ice Builder for Visual Studio.

C++ run-time
The Ice binary distribution for Windows installs the Visual Studio 2013 32-bit C++ run-time on your computer. It does not install any other C++ run-time.

This allows you to execute all C++ programs installed in the main `bin` directory (such as `icegridregistry.exe` and `slice2java.exe`) without having to install Visual Studio 2013 or the Visual Studio 2013 C++ x64 redistributable package. However, if you want to execute C++ programs installed in the bin subdirectories, such as `bin\x64\icegridregistry.exe` or `bin\vc110\icegridregistry.exe`, you need to install the corresponding Visual Studio C++ redistributable package.

Setting up your Windows environment to use Ice

After installing Ice, read the relevant language-specific sections below to learn how to configure your environment and start programming with Ice.

Please note that the instructions below assume that Ice was installed in the default location `C:\Program Files (x86)\ZeroC\Ice-3.6.3`.

C++

To use Ice for C++ you need to add one or more of the Ice `bin` directories to your `PATH`. The changes you must make depend on your compiler and architecture, as listed below. If you don't want to rely on the `PATH` environment variable to locate the Ice for C++ DLLs, you can

Copyright 2016, ZeroC, Inc.

also copy the DLLs into the same directory as your executable.

PATH settings for Visual Studio 2015

Visual Studio 2015, 32-bit:

```
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin\vc140;%PATH%
```

Visual Studio 2015, 64-bit:

```
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin\vc140;%PATH%
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin\vc140\x64;%PATH%
```

The `x64` directory must come first in your `PATH`.

PATH settings for Visual Studio 2013

Visual Studio 2013, 32-bit:

```
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin;%PATH%
```

Visual Studio 2013, 64-bit:

```
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin;%PATH%
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin\x64;%PATH%
```

The `x64` directory must come first in your `PATH`.

PATH settings for Visual Studio 2012

Visual Studio 2012, 32-bit:

```
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin\vc110;%PATH%
```

Visual Studio 2012, 64-bit:

```
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin\vc110;%PATH%
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin\vc110\x64;%PATH%
```

The `x64` directory must come first in your `PATH`.

Compiler Flags

The Release versions of the Ice libraries are compiled with `/MD` to select the multi-threaded Visual C++ run-time library, while the Debug versions use `/MDd` to select the debug multi-threaded run-time library. Both versions of the Ice libraries are compiled with `/EHsc` to select an exception handling model.

Your application must be compiled with the same flags as the Ice libraries:

- Release: `/MD /EHsc`

Copyright 2016, ZeroC, Inc.

- Debug: `/MDd /EHsc`

Import Libraries

You no longer need to list Ice import libraries such as `IceD.lib` when linking with Ice libraries. See Linking with C++ Libraries on Windows for additional details.

.NET

Locating the .NET assemblies

To use Ice for .NET, you can either copy the .NET assemblies to the directory of your executable or add the .NET assemblies to the Global Assembly Cache (GAC).

Copying the Ice for .NET assemblies to the executable directory is the simplest solution. You can set up your Visual Studio projects to copy the assemblies by setting the `Copy Local` property to `True`. To access this property in the Solution Explorer, open the `References` folder of your project and click on the assembly to access its properties in the `Properties` panel.

You can also add the Ice for .NET assemblies to the GAC. To do this, open a command prompt window with Administrator privileges and use `gacutil`:

```
gacutil /i <library.dll>
```

The `gacutil` tool is included with your Visual Studio installation. You can find which `gacutil` you are using with the `where` command, for example:

```
where gacutil
C:\Program Files (x86)\Microsoft SDKs\Windows\v8.1A\bin\NETFX 4.5.1
Tools\x64\gacutil.exe
```

Once installed in the GAC, the assemblies will always be located correctly without having to set environment variables or copy them into the same directory as an executable.

Line numbers for stack traces

If you've installed the Ice assemblies in the GAC and want line numbers for stack traces, you must also install the PDB (.pdb) files in the GAC. Unfortunately, you cannot do this using Explorer, so you have to do it from the command line. Open a command shell window and navigate to `C:\Windows\Microsoft.NET\assembly\GAC_MSIL\Ice` (assuming `C:\Windows` is your system root). Doing a directory listing there, you will find a directory named `v4.0_3.6.3.0__<UUID>`, for example:

`v4.0_3.6.3.0__cdd571ade22f2f16`

Change to that directory (making sure that you use the correct version number for this release of Ice). In this directory, you will see the `Ice.dll` you installed into the GAC in the preceding step. Now copy the `Ice.pdb` file into this directory:

```
copy C:\Program Files (x86)\ZeroC\Ice-3.6.3\Assemblies\Ice.pdb
```

Using protocol compression

The Ice for .NET run time implements protocol compression by dynamically loading the native library `bzip2.dll` from a directory in your `PATH`. Ice disables the protocol compression feature if it is unable to load the `bzip2` library successfully.

This DLL is included in your Ice distribution, therefore the Ice `bin` directory must be added to your `PATH`:

Copyright 2016, ZeroC, Inc.

```
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin;%PATH%
```

On 64-bit Windows, use the following setting instead:

```
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin\x64;%PATH%
```

If the wrong PATH is set, the Ice run time prints a warning to the console when it detects a `bzip2.dll` format mismatch during start-up.

Using managed code in .NET

The main Ice for .NET assembly (`Ice.dll`) included in the Windows binary distribution uses unmanaged code. If you require only managed code then you can clone the ice repository and build Ice for .NET in a purely managed version. Note that the managed version of Ice for .NET omits support for protocol compression and for signal handling in the `Ice.Application` class.

Java

JAR Files

The preferred way to use Ice for Java is to configure your build environment to fetch the build artifacts from the ZeroC Maven repository located at:

> http://repo.zeroc.com/nexus/content/repositories/releases

The Maven package ids are as follows:

- `groupId=com.zeroc, version=3.6.3, artifactId=ice,`
- `groupId=com.zeroc, version=3.6.3, artifactId=freeze`
- `groupId=com.zeroc, version=3.6.3, artifactId=glacier2`
- `groupId=com.zeroc, version=3.6.3, artifactId=icebox`
- `groupId=com.zeroc, version=3.6.3, artifactId=icediscovery`
- `groupId=com.zeroc, version=3.6.3, artifactId=icegrid`
- `groupId=com.zeroc, version=3.6.3, artifactId=icelocatordiscovery`
- `groupId=com.zeroc, version=3.6.3, artifactId=icepatch2`
- `groupId=com.zeroc, version=3.6.3, artifactId=icestorm`

If you cannot fetch the Ice for Java artifacts from the ZeroC Maven repository, the Ice JAR files are also installed locally in `C:\Program Files (x86)\ZeroC\Ice-3.6.3\lib`. The Ice JAR files are named as follows:

- `ice-3.6.3.jar`
- `freeze-3.6.3.jar`
- `glacier2-3.6.3.jar`
- `icebox-3.6.3.jar`
- `icediscovery-3.6.3.jar`
- `icegrid-3.6.3.jar`
- `icelocatordiscovery-3.6.3.jar`
- `icepatch2-3.6.3.jar`
- `icestorm-3.6.3.jar`

Freeze and Berkeley DB

In order to use Freeze for Java, you need to add both `freeze-3.6.3.jar` and the Berkeley DB jar file (`db.jar`) to your build path. `db.jar` is installed alongside the Ice jar files in `C:\Program Files (x86)\ZeroC\Ice-3.6.3\lib`.

At run time, the JVM requires that the directory containing Berkeley DB's native libraries be included in `java.library.path`, which corresponds to the PATH environment variable on Windows.

With a 32-bit JVM, update your PATH as follows:

Copyright 2016, ZeroC, Inc.

```
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin;%PATH%
```

For a 64-bit JVM, use the following setting instead:

```
set PATH=C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin\x64;%PATH%
```

Bzip2

Ice for Java supports protocol compression using the `bzip2` classes included with ant. Compression is automatically enabled if these classes are present in your `CLASSPATH`. You can either add `ant.jar` to your `CLASSPATH`, or download only the bzip2 classes from http://www.kohsuke.org/bzip2.

You can also retrieve this bzip2 jar file from ZeroC's Maven repository located at http://repo.zeroc.com/nexus/content/repositories/thirdparty.

The Maven package id for the bzip2 jar file is as follows:

- `groupId=org.apache.tools, version=1.0, artifactId=bzip2`

These classes are a pure Java implementation of the bzip2 algorithm and therefore add significant latency to Ice requests.

PHP

The Ice extension for PHP included in this installer is compiled with VC11 for compatibility with the PHP binary distribution. To use this extension, you will need a compatible PHP binary distribution as well as a compatible Web server. If you wish to use Apache, you can obtain a VC11 build of Apache from alternate sources as Apache does not provide a VC11 build (PHP recommends the builds by Apache Lounge).

If you require a version of the Ice extension for a different environment, you will need to compile the extension from source. Clone the ice repository and review the `php/INSTALL` file for details.

We provide instructions below for configuring PHP to use the Ice extension. Please defer to your web server's instructions to configure it to run PHP. These instructions make several assumptions:

- PHP is installed in `C:\PHP`
- `C:\PHP` is in your system path
- Ice is installed in `C:\Program Files (x86)\ZeroC\Ice-3.6.3`

If you have a different installation, you will need to make the appropriate changes as you follow the instructions.

1. Determine the location of your PHP configuration file

Check for the location of your `php.ini` file by examining the output of the `php -i` command, or by calling the `phpinfo()` function from a script, for the line `Loaded Configuration File`. For example:

```
C:\PHP\php.ini
```

If no file has been loaded then you will need to create one. PHP provides templates for development and production use:

- `C:\PHP\php.ini-development`
- `C:\PHP\php.ini-production`

Depending on your environment, copy the correct file to a new file, `C:\PHP\php.ini`, and configure it as needed. Verify that it is loaded as described above.

As an administrator, open the `php.ini` in a text editor and append the following line:

```
extension = php_ice.dll
```

The file `php_ice.dll` contains the Ice extension for PHP.

Copyright 2016, ZeroC, Inc.

2. Install the Ice extension DLL

Look for the `extension_dir` setting in the output of the command `php -i`:

```
extension_dir = "C:\PHP\ext"
```

If instead the `extension_dir` setting contains a relative path, it is resolved relative to the working directory of the web server process (For example, Apache's working directory is usually its installation directory.)

Copy the DLL for the Ice extension to PHP's extension directory:

```
copy "C:\Program Files (x86)\ZeroC\Ice-3.6.3\php\php_ice.dll" "C:\PHP\ext"
```

3. Verify that the dependent libraries can be loaded

Regardless of the location of PHP's extension directory, the Ice extension's dependent libraries must be located in the executable search path of the calling process, either your web server or directly from the command line.

The Ice extension depends on the following libraries:

- `bzip2.dll`
- `ice36.dll`
- `iceutil36.dll`
- `slice36.dll`

All of these files can be found in the `bin` subdirectory of your Ice installation (e.g., `C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin\vc110`). To add the Ice installation directory to the system `PATH`, open the System control panel and change the system `PATH` to include `C:\Program Files (x86)\ZeroC\Ice-3.6.3\bin\vc110`. Note that Windows must be restarted for this change to take effect.

If you're using Apache and it cannot find or access a DLL, Apache startup may fail with an access violation, or the PHP module may ignore the Ice extension and continue its initialization. Consequently, a successful Apache startup does not necessarily mean that the Ice extension has been loaded. Unfortunately, the message reported by PHP in Apache's error log is not very helpful; the error might imply that it cannot find `php_ice.dll` when in fact it was able to open `php_ice.dll` but a dependent DLL was missing or inaccessible.

4. Review access rights

Review the access rights on PHP's extension directory, the Ice extension DLL, and its dependent libraries. For example when running as a Windows service, Apache runs in the `Local System` account (also known as `NT Authority\SYSTEM`). You can use the `icacls` utility in a command window to view and modify access rights. For example, run the following commands to review the current access rights of the Ice extension:

```
cd \PHP\ext
icacls php_ice.dll
```

5. Locate the Ice run-time files

Your application will also need to include at least some of the Ice for PHP run-time source files (installed in `C:\Program Files (x86)\ZeroC\Ice-3.6.3\php`). To make these files available to your application, you can either modify PHP's include path or copy the necessary files to a directory that is already in the interpreter's include path. You can determine the current include path by running the command `php -i` and searching for an entry named `include_path`.

If you want to make the Ice run-time files available to all PHP applications on the host, you can modify the `include_path` setting in `php.ini` to add the installation directory:

```
include_path = C:\Program Files (x86)\ZeroC\Ice-3.6.3;...
```

Another option is to modify the include path from within your script prior to including any Ice run-time file:

Copyright 2016, ZeroC, Inc.

PHP

```
ini_set('include_path', ini_get('include_path') . PATH_SEPARATOR . 'C:/Program
Files (x86)/ZeroC/Ice-3.6.3/php')
require 'Ice.php'; // Load the core Ice run time definitions.
```

Using the sample programs on Windows

The Ice sample programs are provided in a GitHub repository. You can browse this repository to see build and usage instructions for all supported programming languages.

During installation, you can choose to install the sample programs for all supported programming languages in a folder of your choice (by default a sub-folder of your `Documents` folder). You'll then find a shortcut to this folder in your Start menu.

Alternatively, you can clone the ice-demos repository and use its 3.6 branch:

```
git clone -b 3.6 https://github.com/zeroc-ice/ice-demos.git
cd ice-demos
```

Configuration files for IceGrid and Glacier2 services

The `config` subdirectory of your Ice installation includes sample configuration files for the Glacier2 router, IceGrid node, and IceGrid registry. These files provide a good starting point on which to base your own configurations, and they contain comments that describe the settings in detail.

If you intend to edit one of the configuration files but you have installed your distribution into `\Program Files` or `\Program Files (x86)`, be aware that Windows makes it difficult to make permanent modifications to files in these directories. We recommend copying the files to the location of your choice. For more information on this topic, refer to the **Virtualization** section of New UAC Technologies for Windows Vista.

The Ice manual provides more information on installing and running the IceGrid registry, IceGrid node, and Glacier2 router as Windows services.

Starting IceGrid Admin on Windows

You can launch IceGrid Admin using the shortcut that the Ice installer created in your Start menu as **IceGrid Admin**. IceGrid Admin is a Java program and requires JRE 7u6 or later.

Unattended installation

The Ice installer supports unattended installation. For example, in an administrative command window you can run:

```
start /wait Ice-3.6.3.msi /qn /l*v install.log
```

Windows may prompt you to confirm the installation, otherwise the installer runs using its default configuration (i.e., default installation directory with no optional components) but without any user interface. The installer will create a log of its activities in the file `install.log`.

Use the `/qb` option instead to install the default configuration while showing a progress dialog:

```
start /wait Ice-3.6.3.msi /qb /l*v install.log
```


Copyright 2016, ZeroC, Inc.

Ice installer registry key

The Ice installer adds information to the Windows registry to indicate where it was installed. Developers can use this information to locate the Ice files in their applications.

The registration key used by this installer is:

```
HKEY_LOCAL_MACHINE\Software\ZeroC\Ice 3.6.3
```

The install location is stored as a string value named `InstallDir`.

Third-party packages for Windows

The binary distribution for Windows includes the following third-party packages as separate binaries:

- Berkeley DB 5.3.28 (C/C++ and Java run-time)
- Bzip2 1.0.6 (C run-time)
- Expat 2.1.0 (C run-time)

Copyright 2016, ZeroC, Inc.

Using the Linux Binary Distributions

This page provides important information for users of the Ice binary distributions on Linux platforms. You can obtain these distributions at the ZeroC web site.

On this page:

- Overview of the Ice Binary Distributions for Linux
  - DEB Packages
  - RPM Packages
  - Bi-Arch Support
- Installing the Linux Distributions
  - Installing Ice on Ubuntu
  - Installing Ice on Red Hat Enterprise Linux 7
  - Installing Ice on Red Hat Enterprise Linux 6
  - Installing Ice on Amazon Linux 2016.09
  - Installing Ice on SuSE Enterprise Server 12
  - Installing Ice on SuSE Enterprise Server 11
- Setting up your Linux environment to use Ice
  - C++
  - Java
  - PHP
- Using the sample programs on Linux
- Starting IceGrid Admin on Linux
- Startup scripts for IceGrid and Glacier2 services

Overview of the Ice Binary Distributions for Linux

DEB Packages

ZeroC provides the following DEB packages for Ubuntu:

Package	Description
zeroc-ice-all-dev	Meta package that installs all development packages
zeroc-ice-all-runtime	Meta package that installs all run-time packages, servers and utilities
libzeroc-freeze3.6	Freeze for C++ run-time libraries
libzeroc-ice-dev	C++ header files, libraries, and Slice compiler for C++
libzeroc-ice-java	Java development kit, including JAR files and Slice compiler for Java
libzeroc-ice3.6	C++ run-time libraries
libzeroc-icestorm3.6	IceStorm service for IceBox C++
php5-zeroc-ice	PHP extension and run time files
php5-zeroc-ice-dev	Slice compiler for PHP development
zeroc-glacier2	Glacier2 service
zeroc-ice-slice	Slice files
zeroc-ice-utils	Utilities necessary for administering an Ice installation
zeroc-ice-utils-java	Additional utilities that require Java, such as the IceGrid Admin GUI
zeroc-icebox	IceBox server for C++
zeroc-icegrid	IceGrid service
zeroc-icepatch2	IcePatch2 service

Copyright 2016, ZeroC, Inc.

RPM Packages

ZeroC provides the following RPMs for Red Hat Enterprise Linux, SuSE Linux Enterprise Server, and Amazon Linux:

RPM	Description
ice-all-devel	Meta package that installs all development packages
ice-all-runtime	Meta package that installs all run-time packages, servers and utilities
glacier2	Glacier2 service
ice-utils	Utilities necessary for administering an Ice installation
ice-utils-java	Additional utilities that require Java, such as the IceGrid Admin GUI
ice-slice	Slice files
icebox	IceBox server for C++
icegrid	IceGrid service
icepatch2	IcePatch2 executable
libfreeze3.6-c++	Freeze for C++ run-time libraries
libice-c++-devel	C++ header files, libraries, and Slice compiler for C++
libice-java	Java development kit, including JAR files and Slice compiler for Java
libice3.6-c++	C++ run-time libraries
libicestorm3.6	IceStorm service
php-ice	PHP extension and run time files
php-ice-devel	Slice compiler for PHP development

ZeroC also supplies RPMs for the following third-party packages:

RPM	Description
db53	Berkeley DB 5.3.28 C and C++ run time libraries
db53-devel	C++ development files for Berkeley DB 5.3.28
db53-java	Berkeley DB 5.3.28 Java run time
db53-utils	Berkeley DB 5.3.28 command-line utilities
mcpp-devel	MCPP C++ preprocessor library

The db53 packages are supplied only for Amazon Linux, Red Hat Enterprise Linux 6 and SuSE Linux Enterprise Server. The db53-devel and the mcpp-devel RPMs are only necessary for building Ice from source.

On Red Hat Enterprise Linux 7, Berkeley DB 5.3 is available as packages libdb-cxx, libdb-java and libdb-utils; some of these packages are in the rhel-7-server-optional-rpms repository that you may need to enable on your system.

Bi-Arch Support

Recent Linux releases such as RHEL 7 are moving away from separate 64-bit and 32-bit distributions and migrating to a single 64-bit release that includes support for 32-bit packages. On these systems, all of the Ice packages listed above are provided for the x86_64

Copyright 2016, ZeroC, Inc.

architecture, along with a limited subset of packages for the x86 architecture. The subset includes C++ run-time and development libraries, along with 32-bit versions of the IceBox server and IceStorm service. The 32-bit IceBox package installs the executables as `icebox32` and `icebox32++11`; the latter is used for IceBox services built in C++11 mode, as described in the C++ section below.

For development purposes, you will still need to install the 64-bit development kit packages: the 32-bit development kit packages complement these 64-bit packages.

On other systems, such as Amazon Linux and SLES 12, we only provide packages for the x86_64 architecture.

Installing the Linux Distributions

This section describes how to install Ice binary packages for all of the supported Linux platforms.

Installing Ice on Ubuntu

Follow the instructions below to install Ice on Ubuntu.

1. Install ZeroC's key to avoid warnings with unsigned packages:

```
sudo apt-key adv --keyserver keyserver.ubuntu.com --recv 5E6DA83306132997
```

2. Add the Ice repository to your system:

```
sudo apt-add-repository "deb http://zeroc.com/download/apt/ubuntu`lsb_release -rs`
stable main"
```

3. Update the package list and install:

```
sudo apt-get update
sudo apt-get install zeroc-ice-all-runtime zeroc-ice-all-dev
```

Refer to the package summary if you would like to install fewer packages.

The metapackages `zeroc-ice-all-runtime` and `zeroc-ice-all-dev` reference the required Berkeley DB dependencies from the Universe repository. You will need to install `db5.3-util` separately to perform admin tasks such as hot backups on your Berkeley DB databases with command-line utilities:

```
sudo apt-get install db5.3-util
```

4. Install the source package (optional):

```
sudo apt-get source zeroc-ice3.6
```

Installing Ice on Red Hat Enterprise Linux 7

Follow the instructions below to install Ice on RHEL 7:

1. Add the Ice repositories to your system:

Copyright 2016, ZeroC, Inc.

```
cd /etc/yum.repos.d
sudo wget https://zeroc.com/download/rpm/zeroc-ice-el7.repo
```

2. Install Ice:

```
sudo yum install ice-all-runtime ice-all-devel
```

Refer to the package summary if you would like to install fewer packages.

The yum documentation provides more information about installing packages on RHEL 7.

Notes:

- Although RHEL 7 is a 64-bit operating system, we continue to support 32-bit applications by providing i686 versions of IceBox, IceStorm, C++ run-time libraries, and C++ development libraries, for example:

  ```
  sudo yum install libzeroc-ice-dev.i686
  ```

- This distribution depends on the Berkeley DB packages included with RHEL 7:

  ```
  sudo yum list 'libdb*'
  ```

  You will need to install `libdb-utils` to perform admin tasks such as hot backups on your Berkeley DB databases with command-line utilities:

  ```
  sudo yum install libdb-utils
  ```

Installing Ice on Red Hat Enterprise Linux 6

Follow the instructions below to install Ice on RHEL 6:

1. Add the Ice repositories to your system:

```
cd /etc/yum.repos.d
sudo wget https://zeroc.com/download/rpm/zeroc-ice-el6.repo
```

2. Install Ice:

```
sudo yum install ice-all-runtime ice-all-devel
```

Refer to the package summary if you would like to install fewer packages.

The yum documentation provides more information about installing packages on RHEL 6.

Notes:

- This distribution depends on Berkeley DB packages included in our third-party repository:

Copyright 2016, ZeroC, Inc.

```
sudo yum list 'db53*'
```

You will need to install `db53-utils` to perform admin tasks such as hot backups on your Berkeley DB databases with command-line utilities:

```
sudo yum install db53-utils
```

Installing Ice on Amazon Linux 2016.09

Follow the instructions below to install Ice on Amazon Linux:

1. Add the Ice repositories to your system:

```
cd /etc/yum.repos.d
sudo wget https://zeroc.com/download/rpm/zeroc-ice-amzn1.repo
```

2. Install Ice:

```
sudo yum install ice-all-runtime ice-all-devel
```

Refer to the package summary if you would like to install fewer packages.

The EC2 documentation provides more information about installing packages on Amazon Linux.

Notes:

- This distribution depends on Berkeley DB packages included in our third-party repository:

```
sudo yum list 'db53*'
```

You will need to install `db53-utils` to perform admin tasks such as hot backups on your Berkeley DB databases with command-line utilities:

```
sudo yum install db53-utils
```

Installing Ice on SuSE Enterprise Server 12

Follow the instructions below to install Ice on SLES 12:

1. Add the Ice repositories to your system:

```
sudo rpmkeys --import https://zeroc.com/download/GPG-KEY-zeroc-release
wget https://zeroc.com/download/rpm/zeroc-ice-sles12.repo
sudo zypper addrepo zeroc-ice-sles12.repo
```

2. Install Ice:

Copyright 2016, ZeroC, Inc.

```
sudo zypper install ice-all-runtime ice-all-devel
```

Refer to the package summary if you would like to install fewer packages.

Notes:

- This distribution depends on Berkeley DB packages included in our third-party repository:

  ```
  sudo zypper search 'db53*'
  ```

  You will need to install `db53-utils` to perform admin tasks such as hot backups on your Berkeley DB databases with command-line utilities:

  ```
  sudo zypper install db53-utils
  ```

Installing Ice on SuSE Enterprise Server 11

Follow the instructions below to install Ice on SLES 11:

1. Add the Ice repositories to your system:

```
wget https://zeroc.com/download/rpm/zeroc-ice-sles11.repo
sudo zypper addrepo zeroc-ice-sles11.repo
```

2. Install Ice:

```
sudo zypper install ice-all-runtime ice-all-devel
```

Refer to the package summary if you would like to install fewer packages.

Notes:

- This distribution depends on Berkeley DB packages included in our third-party repository:

  ```
  sudo zypper search 'db53*'
  ```

  You will need to install `db53-utils` to perform admin tasks such as hot backups on your Berkeley DB databases with command-line utilities:

  ```
  sudo zypper install db53-utils
  ```

Setting up your Linux environment to use Ice

After installing Ice, read the relevant language-specific sections below to learn how to configure your environment and start programming with Ice.

Copyright 2016, ZeroC, Inc.

C++

A C++ program needs to link with at least `libIce` and `libIceUtil`, so a typical link command would look like this:

```
c++ -o myprogram myprogram.o -pthread -lIce -lIceUtil
```

Additional libraries are necessary if you are using an Ice service such as IceGrid or Glacier2.

Always use the `-pthread` option when compiling and linking your Ice applications.

If you build your C++ program in C++11 or C++14 mode (with `-std=c++0x`, `-std=gnu++0x`, `-std=c++11`, `-std=gnu++11`, `-std=c++14`, `-std=gnu++14`), you need to link with the C++11 libraries included in your binary distribution. For 32-bit programs:

```
c++ -std=c++11 -o myprogram myprogram.o -pthread -L/usr/lib/c++11 -lIce -lIceUtil
```

For 64-bit programs:

```
c++ -std=c++11 -o myprogram myprogram.o -pthread -L/usr/lib64/c++11 -lIce -lIceUtil
```

Java

The Ice for Java package (`libice-java` or `libzeroc-ice-java` depending on your Linux distribution) is a development kit. A production-quality Java application that uses Ice for Java is expected to incorporate all the Ice JAR files it needs; it should not depend on the Ice for Java package.

Gradle Builder

Developers can use the Ice Builder for Gradle to simplify their Java projects. The builder is not included in the binary distributions and must be installed separately.

JAR Files

The preferred way to develop with Ice for Java is to configure your build environment to fetch the build artifacts from the ZeroC Maven repository located at:

http://repo.zeroc.com/nexus/content/repositories/releases

The Maven package ids are as follows:

- `groupId=com.zeroc, version=3.6.3, artifactId=ice,`
- `groupId=com.zeroc, version=3.6.3, artifactId=freeze`
- `groupId=com.zeroc, version=3.6.3, artifactId=glacier2`
- `groupId=com.zeroc, version=3.6.3, artifactId=icebox`
- `groupId=com.zeroc, version=3.6.3, artifactId=icediscovery`
- `groupId=com.zeroc, version=3.6.3, artifactId=icegrid`
- `groupId=com.zeroc, version=3.6.3, artifactId=icelocatordiscovery`
- `groupId=com.zeroc, version=3.6.3, artifactId=icepatch2`
- `groupId=com.zeroc, version=3.6.3, artifactId=icestorm`

If you cannot fetch the Ice for Java artifacts from the ZeroC Maven repository, the Ice JAR files are also installed locally in `/usr/share/java`. The files are named as follows:

- `ice-3.6.3.jar`
- `freeze-3.6.3.jar`
- `glacier2-3.6.3.jar`

Copyright 2016, ZeroC, Inc.

- icebox-3.6.3.jar
- icediscovery-3.6.3.jar
- icegrid-3.6.3.jar
- icelocatordiscovery-3.6.3.jar
- icepatch2-3.6.3.jar
- icestorm-3.6.3.jar

Freeze and Berkeley DB

In order to use Freeze for Java, you need to add both `freeze-3.6.3.jar` and the Berkeley DB JAR file (`db.jar`) to your build path. `db.jar` is installed alongside the Ice JAR files in `/usr/share/java`.

To use Freeze, the JVM also requires that the directory containing Berkeley DB's native libraries be listed in `java.library.path`, therefore you must add this directory to your `LD_LIBRARY_PATH`. Assuming you are using ZeroC's distribution of Berkeley DB, the bash command is shown below:

```
export LD_LIBRARY_PATH=/usr/lib:$LD_LIBRARY_PATH                (RHEL, SLES,
Amazon)
export LD_LIBRARY_PATH=/usr/lib/i386-linux-gnu:$LD_LIBRARY_PATH  (Ubuntu)
```

On an x86_64 system with a 64-bit JVM, the 64-bit Berkeley DB libraries are installed in a different directory:

```
export LD_LIBRARY_PATH=/usr/lib64:$LD_LIBRARY_PATH                (RHEL, SLES,
Amazon)
export LD_LIBRARY_PATH=/usr/lib/x86_64-linux-gnu:$LD_LIBRARY_PATH  (Ubuntu)
```

Bzip2

Ice for Java supports protocol compression using the `bzip2` classes included with ant. Compression is automatically enabled if these classes are present in your `CLASSPATH`. You can either add `ant.jar` to your `CLASSPATH`, or download only the bzip2 classes from http://www.kohsuke.org/bzip2.

You can also retrieve this bzip2 JAR file from ZeroC's Maven repository located at http://repo.zeroc.com/nexus/content/repositories/thirdparty .

The Maven package id for the bzip2 JAR file is as follows:

- `groupId=org.apache.tools, version=1.0, artifactId=bzip2`

These classes are a pure Java implementation of the bzip2 algorithm and therefore add significant latency to Ice requests.

PHP

The Ice extension for PHP is loaded automatically when the interpreter loads the contents of the file `/etc/php.d/ice.ini` (on Red Hat Enterprise Linux and Amazon Linux) or `/etc/php5/conf.d/ice.ini` (on SuSE Linux Enterprise Server and Ubuntu). This file contains the line shown below:

```
extension=IcePHP.so
```

You can modify this file to include additional configuration directives.

At run time, the PHP interpreter requires the Ice shared libraries.

You can verify that the Ice extension is installed properly by examining the output of the `php -m` command, or by calling the `phpinfo()` fun

Copyright 2016, ZeroC, Inc.

ction from a script.

Your application will also need to include at least some of the Ice for PHP run-time source files (installed in `/usr/share/php` on RHEL, Amazon Linux, and Ubuntu, and in `/usr/share/php5` on SLES). This installation directory is included in PHP's default include path, which you can verify by executing the following command:

```
php -i | grep include_path
```

If the installation directory is listed, no further action is necessary to make the run-time source files available to your application. Otherwise, you can modify the `include_path` setting in `php.ini` to add the installation directory:

```
include_path = /usr/share/php:...
```

Another option is to modify the include path from within your script prior to including any Ice run-time file:

PHP

```
ini_set('include_path', ini_get('include_path') . PATH_SEPARATOR .
'/usr/share/php');
require 'Ice.php'; // Load the core Ice run time definitions.
```

SELinux Notes (for Red Hat Enterprise Linux users)

SELinux augments the traditional Unix permissions with a number of new features. In particular, SELinux can prevent the `httpd` daemon from opening network connections and reading files without the proper SELinux types.

If you suspect that your PHP application does not work due to SELinux restrictions, we recommend that you first try it with SELinux disabled. As root, run:

```
setenforce 0
```

to disable SELinux until the next reboot of your computer.

If you want to run `httpd` with the Ice extension and SELinux enabled, you must do the following:

1. Allow httpd to open network connections:

   ```
   setsebool httpd_can_network_connect=1
   ```

   You can add the -P option to make this setting persistent across reboots.

2. Make sure any `.ice` file used by your PHP scripts can be read by `httpd`. The enclosing directory also needs to be accessible. For example:

   ```
   chcon -R -t httpd_sys_content_t /opt/MyApp/slice
   ```

For more information on SELinux in Red Hat Enterprise Linux, refer to this Red Hat document.

Copyright 2016, ZeroC, Inc.

Using the sample programs on Linux

Sample programs for all programming languages are available in a separate GitHub repository. Simply clone this repository and use its 3.6 branch:

```
git clone -b 3.6 https://github.com/zeroc-ice/ice-demos.git
cd ice-demos
```

Starting IceGrid Admin on Linux

You can launch IceGrid Admin with the **`icegridgui`** command. IceGrid Admin is a Java program and requires JRE 7u6 or later.

Startup scripts for IceGrid and Glacier2 services

All distributions include the following service configuration files:

- `/etc/icegridregistry.conf`
- `/etc/icegridnode.conf`
- `/etc/glacier2router.conf`

Distributions which use `sysvinit` or `upstart` contain the following sample scripts:

- `/etc/init.d/icegridregistry`
- `/etc/init.d/icegridnode`
- `/etc/init.d/glacier2router`

Distributions which use `systemd` contain the following services:

- `icegridregistry.service`
- `icegridnode.service`
- `glacier2router.service`

The installation also creates a user account and group for running these services (account `ice` and group `ice`), and data directories for `ice gridregistry` and `icegridnode` (`/var/lib/ice/icegrid/registry` and `/var/lib/ice/icegrid/node1`).

By default, all these services are off at all runlevels. You need to manually switch on one or more runlevels, for example:

Copyright 2016, ZeroC, Inc.

```
#
# On a Red Hat Enterprise Linux 6 system, configure the icegridregistry
# to start at runlevel 3 and 5:
#
sudo chkconfig --level 35 icegridregistry on

#
# On a SuSE Linux Enterprise Server 11 system, configure the
# icegridregistry to start at runlevel 3 and 5:
#
sudo chkconfig icegridregistry 35

#
# On an Ubuntu system, configure the icegridregistry to start at the
# default run levels:
#
sudo update-rc.d icegridregistry defaults

#
# On systems using systemd, start icegridregistry with the multi-user target
#
sudo systemctl enable icegridregistry.service
```

Before doing so, please review the script itself and its associated configuration file.

Copyright 2016, ZeroC, Inc.

Using the macOS Binary Distribution

This page provides important information for users of the Ice binary distribution for macOS.

On this page:

- Overview of the macOS binary distribution
- Installing the macOS binary distribution
- Setting up your macOS environment to use Ice
  - C++
  - Objective-C
  - PHP
  - Java
  - Berkeley DB utilities
- Using the sample programs on macOS
- Starting IceGrid Admin on macOS

Overview of the macOS binary distribution

Ice for C++, Java, Objective-C, and PHP on macOS are provided through Homebrew. The `ice` formula includes the following components by default:

- Run-time libraries for C++, Objective-C and PHP
- Run-time libraries for the Freeze service in C++
- Executables for IceGrid, IceStorm, Glacier2, and IcePatch2 services
- Tools and libraries for developing Ice applications

More information about Homebrew is available at http://brew.sh.

By default, the `ice` formula does build any Java component - this includes Ice for Java, the Freeze service for Java, and the IceGrid Admin graphical tool.

Installing the macOS binary distribution

Using Homebrew, you can install the distribution with this command:

```
brew install ice [--with-java]
```

The `--with-java` option builds the Java components and the IceGrid Admin app. You can also install IceGrid Admin on its own by downloading IceGrid Admin.dmg.

Alternatively, you can install Ice from ZeroC's Homebrew tap.

Setting up your macOS environment to use Ice

After installing Ice, read the relevant language-specific sections below to learn how to configure your environment and start programming with Ice.

C++

Xcode Builder

We recommend installing the Ice Builder for Xcode, which automates the compilation of Slice files to C++ and Objective-C.

Compiling and Linking

The distribution includes libraries built with LLVM `libc++`. These libraries also support C++11 applications when compiled with the `--std=c++11` option.

Copyright 2016, ZeroC, Inc.

When compiling Ice for C++ programs, you must pass the `-pthread` option. A typical compile command would look like this:

```
c++ -c -pthread myprogram.cpp
```

When linking a program you must link with at least `libIce` and `libIceUtil`. A typical link command would look like this:

```
c++ -o myprogram myprogram.o -lIce -lIceUtil
```

Additional libraries are necessary if you are using an Ice service such as IceGrid or Glacier2.

If you want to include Ice in your application bundle, you will need to copy the necessary Ice libraries to the `Contents/Frameworks` subdirectory of your bundle and use `@loader_path/../Frameworks` as the run path when linking the application.

Please refer to the `dyld` man page on your macOS system to learn more about `@loader_path`.

Objective-C

Xcode Builder

We recommend installing the Ice Builder for Xcode, which automates the compilation of Slice files to C++ and Objective-C.

Compiling and Linking

When compiling Ice for Objective-C programs, you must pass the `-pthread` option. A typical compile command would look like this:

```
cc -c -pthread myprogram.m
```

When linking a program you must link with `libIceObjC`. A typical link command would look like this:

```
cc -o myprogram myprogram.o -lIceObjC -framework Foundation
```

Additional libraries are necessary if you are using an Ice service such as IceGrid or Glacier2.

If you want to include Ice in your application bundle, you will need to copy the necessary Ice libraries to the `Contents/Frameworks` subdirectory of your bundle and use `@loader_path/../Frameworks` as the run path when linking the application.

Please refer to the `dyld` man page on your macOS system to learn more about `@loader_path`.

PHP

The macOS distribution includes the Ice extension for PHP and the Slice-to-PHP compiler (`slice2php`).

The PHP interpreter loads the Ice extension automatically when you add the following directive to `/etc/php.ini`:

```
extension=/usr/local/php/extensions/IcePHP.dy
```

Additional configuration directives can also be added to `/etc/php.ini`.

You can verify that the Ice extension is installed properly by examining the output of the `php -m` command, or by calling the `phpinfo()` function from a script.

Your application will also need to include at least some of the Ice for PHP run-time source files (installed in `/usr/local/lib/share/php`)

Copyright 2016, ZeroC, Inc.

. You can modify the `include_path` setting in `php.ini` to add the installation directory:

```
include_path = /usr/local/lib/share/php:...
```

Another option is to modify the include path from within your script prior to including any Ice run-time file:

PHP

```
ini_set('include_path', ini_get('include_path') . PATH_SEPARATOR .
'/usr/local/opt/ice/lib/share/php');
require 'Ice.php'; // Load the core Ice run time definitions.
```

Java

Gradle Builder

Developers can use the Ice Builder for Gradle to simplify their Java projects. The builder is not included in the distribution and must be installed separately.

JAR Files

The preferred way to use Ice for Java is to configure your build environment to fetch the build artifacts from the ZeroC Maven repository located at:

http://repo.zeroc.com/nexus/content/repositories/releases

You can create Java apps with Ice for Java on macOS without installing Ice for Java. The Slice-to-Java (`slice2java`) compiler is always included in your Ice installation, even if you did not install Ice using the `--with-java` option.

The `--with-java` option is necessary only if you develop with Freeze for Java.

The Maven package ids are as follows:

- `groupId=com.zeroc, version=3.6.3, artifactId=ice,`
- `groupId=com.zeroc, version=3.6.3, artifactId=freeze`
- `groupId=com.zeroc, version=3.6.3, artifactId=glacier2`
- `groupId=com.zeroc, version=3.6.3, artifactId=icebox`
- `groupId=com.zeroc, version=3.6.3, artifactId=icediscovery`
- `groupId=com.zeroc, version=3.6.3, artifactId=icegrid`
- `groupId=com.zeroc, version=3.6.3, artifactId=icelocatordiscovery`
- `groupId=com.zeroc, version=3.6.3, artifactId=icepatch2`
- `groupId=com.zeroc, version=3.6.3, artifactId=icestorm`

If you cannot fetch the Ice for Java artifacts from the ZeroC Maven repository, the Ice JAR files are installed locally in `/usr/local/share/java` when you install Ice using the `--with-java` option. The JAR files are named as follows:

- `ice-3.6.3.jar`
- `freeze-3.6.3.jar`
- `glacier2-3.6.3.jar`
- `icebox-3.6.3.jar`
- `icediscovery-3.6.3.jar`
- `icegrid-3.6.3.jar`
- `icelocatordiscovery-3.6.3.jar`
- `icepatch2-3.6.3.jar`
- `icestorm-3.6.3.jar`

Freeze and Berkeley DB

Copyright 2016, ZeroC, Inc.

In order to use Freeze for Java, you need to add both `freeze-3.6.3.jar` and the Berkeley DB JAR file (`db.jar`) to your build path.

ZeroC's Maven Repository does not provide `db.jar`, as this JAR file is a wrapper for a C library. Both `db.jar` and the associated C library are built and installed when you install Ice from Homebrew with the `--with-java` option.

Bzip2

Ice for Java supports protocol compression using the `bzip2` classes included with ant. Compression is automatically enabled if these classes are present in your `CLASSPATH`. You can either add `ant.jar` to your `CLASSPATH`, or download only the bzip2 classes from http://www.kohsuke.org/bzip2.

You can also retrieve this bzip2 JAR file from ZeroC's Maven repository located at http://repo.zeroc.com/nexus/content/repositories/thirdparty

The Maven package id for the bzip2 JAR file is as follows:

- `groupId=org.apache.tools, version=1.0, artifactId=bzip2`

These classes are a pure Java implementation of the bzip2 algorithm and therefore add significant latency to Ice requests.

Berkeley DB utilities

Berkeley DB utilities are installed in`/usr/local/opt/ice/libexec/bin`. You may add this directory to your `PATH`, or call the executables directly.

Using the sample programs on macOS

Sample programs for all programming languages are available in a separate GitHub repository. Simply clone this repository and use its 3.6 branch:

```
git clone -b 3.6 https://github.com/zeroc-ice/ice-demos.git
cd ice-demos
```

Starting IceGrid Admin on macOS

You can launch IceGrid Admin with the `IceGridAdmin` application installed in your `/Applications` directory. This app is a Java program and requires JRE 7u6 or later.

Copyright 2016, ZeroC, Inc.

Using Ice on Android

This page provides important information for users of the Ice for Java distribution.

On this page:

- Overview of Ice on Android
- Configuring an Android Studio project
- Using the sample programs

Overview of Ice on Android

Ice is available for Android though the Ice for Java distribution. To develop an Ice application for Android you will need the following components:

- Binary distribution of Ice for your development platform (Windows, macOS, Linux)
- Android Studio
- Ice Builder for Gradle

Configuring an Android Studio project

Android Studio uses a Gradle-based build system. Using Ice requires modifying your project's `build.gradle` settings file(s) to include the Ice jars. For convenience we recommend using the Ice Builder for Gradle plugin to help automate the compilation of your Slice files.

The following snippets give an example of configuring the Ice and Ice Builder for Gradle sections for an Android project.

Copyright 2016, ZeroC, Inc.

Ice build.gradle configuration

```
buildscript {
    repositories {
        maven {
            url 'https://repo.zeroc.com/nexus/content/repositories/releases'
        }
    }
    dependencies {
        classpath group: 'com.zeroc.gradle.ice-builder', name: 'slice', version:
'1.3.14'
    }
}
apply plugin: 'java'
apply plugin: 'slice'
slice {
    java {
        srcDir = '.'
    }
}

repositories {
    maven {
        url 'https://repo.zeroc.com/nexus/content/repositories/releases'
    }
}
dependencies {
    compile 'com.zeroc:ice:3.6.3'
}
```

Using the sample programs

Sample programs are available at our GitHub repository. You can browse this repository to see build and usage instructions for all supported programming languages. Simply clone this repository and use its 3.6 branch:

```
git clone -b 3.6 https://github.com/zeroc-ice/ice-demos.git
cd ice-demos
```

You can then import the directory `java/Android` into Android Studio.

Copyright 2016, ZeroC, Inc.

Using the JavaScript Distribution

This page provides important information for users of the Ice for JavaScript distribution.

On this page:

- Overview of the Ice for JavaScript Distribution
- Installing the Ice for JavaScript Distribution
- Using Ice for JavaScript
- Using the sample programs

Overview of the Ice for JavaScript Distribution

Ice for JavaScript is available as npm and bower packages. The following npm packages are provided for Node.js developers:

- ice - Ice for JavaScript run time
- slice2js - Slice-to-JavaScript compiler
- gulp-ice-builder - Gulp plug-in to integrate `slice2js` with the gulp build system

The following bower package is provided for developers of browser-based applications:

- bower-ice - Ice for JavaScript run time

Installing the Ice for JavaScript Distribution

Node.js Applications

Install the Ice for JavaScript run time libraries using `npm`:

```
npm install ice --save
```

This package also includes the browser version of Ice for JavaScript in `node_modules/lib`.

You can also run `npm install ice@latest --save` to get the latest patches and bug fixes.

Browser Applications

To install with `bower`:

```
bower install ice --save
```

You can also run `bower install ice#master --save` to get the latest patches and bug fixes.

Slice-to-JavaScript Compiler

The Slice-to-JavaScript compiler (`slice2js`) can be installed with this command:

```
npm install slice2js --save-dev
```

This command builds the compiler from source therefore you must have a supported C++ compiler installed.

Copyright 2016, ZeroC, Inc.

In order to use `slice2js` on the command line, use the `npm --global` option to install this package.

The `slice2js` package includes all of the standard Slice files and automatically adds the slice directory to the include file search path.

Using Ice for JavaScript

The Ice manual provides a complete description of the JavaScript mapping.

Node.js Applications

You can use several top-level Ice package in your Node.js applications, as shown below:

JavaScript

```
var Ice      = require('ice').Ice;
var Glacier2 = require('ice').Glacier2;
var IceStorm = require('ice').IceStorm;
var IceGrid  = require('ice').IceGrid;

var communicator = Ice.initialize(process.argv);
var proxy = communicator.stringToProxy("hello:tcp -h localhost -p 10000");
```

Browser Applications

Add the necessary `<script>` tags to your html to include the Ice for JavaScript components you require.

```
<script src="/bower_components/ice/lib/Ice.js"></script>
<script src="/bower_components/ice/lib/Glacier2.js"></script>
<script src="/bower_components/ice/lib/IceStorm.js"></script>
<script src="/bower_components/ice/lib/IceGrid.js"></script>
<script type="text/javascript">


var communicator = Ice.initialize();
var proxy = communicator.stringToProxy("hello:ws -h localhost -p 10002");
</script>
```

Minified versions are available with the `.min.js` extension.

CDN

The Ice for JavaScript browser libraries are also available through cdnjs.

- https://cdnjs.com/libraries/ice

Copyright 2016, ZeroC, Inc.

```
<script src="https://cdnjs.cloudflare.com/ajax/libs/ice/3.6.3/Ice.min.js"></script>
<script
src="https://cdnjs.cloudflare.com/ajax/libs/ice/3.6.3/Glacier2.min.js"></script>
<script
src="https://cdnjs.cloudflare.com/ajax/libs/ice/3.6.3/IceStorm.min.js"></script>
<script
src="https://cdnjs.cloudflare.com/ajax/libs/ice/3.6.3/IceGrid.min.js"></script>
```

A CDN is a great option for applications that do not wish to host all of their website's files themselves. Among other things, using a CDN will help to reduce bandwidth and allow caching across multiple websites.

Note however that you will still need to serve your client JavaScript files yourself.

Slice-to-JavaScript Compiler

You can execute the Slice-to-JavaScript compiler from code as shown below:

JavaScript

```
var slice2js = require('slice2js');
slice2js.compile(["Hello.ice"]);
```

If you installed the `slice2js` package globally, you can also run `slice2js` on the command line.

Refer to the manual for a description of the arguments accepted by `slice2js`.

Using the sample programs

The Ice sample programs are provided in a GitHub repository. You can browse this repository to see build and usage instructions for all supported programming languages. Simply clone this repository and use its 3.6 branch:

```
git clone -b 3.6 https://github.com/zeroc-ice/ice-demos.git
cd ice-demos
```


Copyright 2016, ZeroC, Inc.

Using the Python Distribution

This page provides important information for users of the Ice for Python distribution.

On this page:

- Overview of the Ice for Python Distribution
- Installing the Ice for Python Distribution
- Using Ice for Python
- Using the sample programs

Overview of the Ice for Python Distribution

Ice for Python is available as a collection of Python packages on the Python Package Index. Each package provides the following components:

- Ice extension for Python
- Standard Slice files
- Slice-to-Python compiler (`slice2py`)

Installing the Ice for Python Distribution

Install Ice for Python using this command:

```
pip install zeroc-ice
```

On Linux this command will build and install the source package. On OS X and Windows it installs a pre-built wheel for the following configurations:

- Windows 32-bit & 64-bit, Python 2.7
- Windows 32-bit & 64-bit, Python 3.5
- macOS 64-bit, Python 2.7

Using Ice for Python

The installation process automatically adds the Ice modules to Python's package directory and adds the Slice-to-Python compiler (`slice2py`) to a directory that's likely already in your executable search path. To verify that Ice is installed, execute these commands:

```
python
>>> import Ice
>>> Ice.getSliceDir()
```

The output of `getSliceDir` shows where the standard Slice files have been installed.

The Ice manual provides a complete description of the Python mapping, including the options for generating Python code from Slice definitions.

Using the sample programs

Sample programs are provided in a separate GitHub repository. You can browse this repository to see build and usage instructions for all supported programming languages. Simply clone this repository and use its 3.6 branch:

Copyright 2016, ZeroC, Inc.

```
git clone -b 3.6 https://github.com/zeroc-ice/ice-demos.git
cd ice-demos
```


Copyright 2016, ZeroC, Inc.

Using the Ruby Distribution

This page provides important information for users of the Ice for Ruby distribution.

On this page:

- Overview of the Ice for Ruby Distribution
- Installing the Ice for Ruby Distribution
- Using Ice for Ruby
- Using the sample programs

Overview of the Ice for Ruby Distribution

Ice for Ruby is available as a collection of Ruby gems. Each package provides the following components:

- Ice extension for Ruby
- Standard Slice files
- Slice-to-Ruby compiler (`slice2rb`)

Installing the Ice for Ruby Distribution

Linux and OS X users can install Ice for Ruby using this command:

```
sudo gem install zeroc-ice
```

This gem builds Ice for Ruby from source and supports Ruby 1.8 or later.

Windows users can install a binary gem instead:

```
gem install zeroc-ice-x64-mingw    (Windows 64-bit)
gem install zeroc-ice-x86-mingw    (Windows 32-bit)
```

These gems require Ruby 2.3.

Using Ice for Ruby

The installation process automatically adds the Ice modules to Ruby's package directory and adds the Slice-to-Ruby compiler (`slice2rb`) to a directory that's likely already in your executable search path. To verify that Ice is installed, execute these commands:

```
irb
irb> require 'Ice'
irb> Ice::getSliceDir()
```

The output of `getSliceDir` shows where the standard Slice files have been installed.

Ruby 1.8
Users of Ruby 1.8 must add the `-rubygems` option when starting the interpreter with `irb` or `ruby`:

Copyright 2016, ZeroC, Inc.

```
irb -rubygems
irb> require 'Ice'
irb> Ice::getSliceDir()
```

The Ice manual provides a complete description of the Ruby mapping, including the options for generating Ruby code from Slice definitions.

Using the sample programs

Sample programs are provided in a separate GitHub repository. You can browse this repository to see build and usage instructions for all supported programming languages. Simply clone this repository and use its 3.6 branch:

```
git clone -b 3.6 https://github.com/zeroc-ice/ice-demos.git
cd ice-demos
```


Copyright 2016, ZeroC, Inc.

Getting Started with Ice on AWS

This page explains how to get started with Ice on the Amazon Web Services (AWS) platform.

- Select and Launch AMI
- Login
- Package Configuration
  - Red Hat Enterprise Linux
  - Ubuntu

Select and Launch AMI

Ice 3.6 is available for AWS on the following AMIs

- **ZeroC Ice for Red Hat Enterprise Linux** - Red Hat Enterprise Linux 7.2
- **ZeroC Ice for Ubuntu** - Ubuntu 14.04 (LTS)

Navigate to the `EC2 Management Console` and follow these steps:

1. Select `Launch Instance`
2. Navigate to the AWS Marketplace page for the ZeroC Ice AMI you wish to use.
3. Choose the instance type you wish to run. ZeroC Ice AMIs are available are available in many sizes.
4. Configure your `Security Group` to allow access to port 22 for SSH login, as well as any ports for any Ice servers you plan on running. Glacier2 listens (typically) on port `4064` (ssl/tls). If you plan on using Glacier then you need to add these port(s) as well.
5. Select which `Key Pair` to use for access to the instance once it is running.
6. Configure any remaining instance settings as desired.
7. Click Accept Terms & Launch with 1-Click. Be sure to read the End-User License Agreement (EULA) beforehand.

Login

You can log in to your instance through SSH.

```
ssh -i <keypair> -l <username> <ip address>
```

Where,

- `<keypair>` - The SSH keypair selected when you launched your instance
- `<username>` - `ec2-user` for Red Hat based instance and `ubuntu` for Ubuntu based instance
- `<ip address>` - IP address of your instance

Package Configuration

ZeroC Ice AMIs have Ice 3.6 package repositories pre-configured. These repositories contain packages for both development and runtime.

Red Hat Enterprise Linux

To install all development and runtime packages:

```
sudo yum install ice-all-runtime ice-all-devel
```

See Using the Linux Binary Distributions for a list of packages.

Ubuntu

To install all development and runtime packages:

Copyright 2016, ZeroC, Inc.

```
sudo apt-get install zeroc-ice-all-runtime zeroc-ice-all-dev
```

See Using the Linux Binary Distributions for a list of packages.

Copyright 2016, ZeroC, Inc.

Building Third Party Packages from Source

Ice has dependencies on a few third-party packages, including Berkeley DB 5.3 and mcpp 2.7.2 (with patches).

The following pages provide build instructions for the third-party packages built and distributed by ZeroC.

Topics

- Building Berkeley DB, bzip2 and mcpp on Windows
- Building Berkeley DB and mcpp on Linux
- Building Berkeley DB on macOS

Copyright 2016, ZeroC, Inc.

Building Berkeley DB, bzip2 and mcpp on Windows

This page describes how to build Berkeley DB, bzip2 and mcpp on Windows for use with Ice.

On this page:

- Berkeley DB
  - Source Archive
  - Patch for Debug DLL
  - Patch for Java
  - Build
- bzip2
  - Source Archive
  - Patch
  - Build
- mcpp
  - Git Repository
  - Build

Berkeley DB

Source Archive

Download the Berkeley DB 5.3.28 source archive from Oracle or ZeroC. You may use the distribution with encryption or without encryption. Since Ice does not use Berkeley DB's encryption features, we recommend downloading the no-encryption (-NC) source archive: https://zeroc.com/download/berkeley-db/db-5.3.28.NC.zip

Then unpack this archive.

Patch for Debug DLL

When building the debug version of the Berkeley DB DLL, you should remove the `DIAGNOSTIC` and `CONFIG_TEST` defines and the `/export:__db_assert` linker option from the project files in `build_windows\VS10`. Without these modifications, database environments created by the debug DLL are not compatible with environments created by the release DLL.

Patch for Java

The Oracle Java 7 JDK and Oracle Java 8 JDK were built with Visual Studio 2010 and we recommend you use the same compiler to build Berkeley DB for Java.

We recommend adding a `_vc100` suffix to the Berkeley DB DLL built with Visual Studio 2010: open the `Berkeley_DB_vs2010` solution and in the `db` project properties, update `Target Name` and add the `_vc100` suffix. Then, to prevent `libdb53_vc100.dll` from having an unnecessary dependency on the Visual Studio C++ run time, remove all the `.cpp` files from the `db` project.

Finally, edit `db_java.vcxproj` and replace `-target 1.5` by `-target 1.7` to use the Java 7 source format.

Build

Follow the Oracle instructions in the `docs/installation/build_win.html` file included in the source distribution.

bzip2

Source Archive

Download the bzip2 1.0.6 source archive from bzip.org or ZeroC.

Then unpack this tar.gz archive.

Patch

Download the following patch: https://zeroc.com/download/bzip2/patch.bzip2-1.0.6

This patch adds support for building bzip2 as a DLL on Windows.

Copyright 2016, ZeroC, Inc.

Apply this patch from the main directory of the bzip2 source distribution:

```
patch -p0 < C:\users\%username%\Downloads\patch.bzip2-1.0.6
```

You can download a `patch` utility for Windows from http://gnuwin32.sourceforge.net/packages/patch.htm

On Windows 7 or later, UAC can make it difficult to use the patch utility unless you take extra steps. One solution is to run `patch.exe` in a command window that you started with Administrator privileges (right-click on Command Prompt in the Start menu and choose "Run as administrator"). If running as administrator is not an option, do the following:

1. Do not install `patch.exe` in a system-protected directory such as `C:\Program Files`.

2. Update the manifest in `patch.exe` as described at http://math.nist.gov/oommf/software-patchsets/patch_on_Windows7.html

Build

Open a Visual Studio command prompt.

If you are using Visual Studio 2010, you can optionally set the `CPP_COMPILER` environment variable to `VC100` to create a DLL with the `_vc100` suffix:

```
set CPP_COMPILER=VC100
```

Change to the main directory of the bzip2 source distribution and type:

```
nmake /f Makefile.mak
```

mcpp

Git Repository

Clone the `zeroc-inc/mcpp` repository hosted on `GitHub`:

```
git clone https://github.com/zeroc-ice/mcpp.git
```

Build

Follow the instructions in README.md.

Copyright 2016, ZeroC, Inc.

Building Berkeley DB and mcpp on Linux

This page describes how to build Berkeley DB and mcpp on Linux, for use with Ice.

On this page:

- Berkeley DB
  - Source Archive
  - Patch
  - Configuration
  - Build and Installation
- mcpp
  - Git Repository
  - Build and Installation

Berkeley DB

Source Archive

Download the Berkeley DB 5.3.28 source archive from Oracle or ZeroC. You may use the distribution with encryption or without encryption. Since Ice does not use Berkeley DB's encryption features, we recommend downloading the no-encryption (-NC) source archive:

```
wget https://zeroc.com/download/berkeley-db/db-5.3.28.NC.tar.gz
```

Then unpack this archive:

```
tar xzf db-5.3.28.NC.tar.gz
```

Patch

Download the following patch: https://zeroc.com/download/berkeley-db/berkeley-db.5.3.28.patch

Apply this patch from the main directory of the berkeley-db source distribution:

```
cd db-5.3.28.NC
patch -p0 < <path to patch>/berkeley-db.5.3.28.patch
```

Configuration

Berkeley DB must be configured with C++ support enabled. If you intend to use Ice for Java with Berkeley DB, you must also enable Java support. For example:

```
cd db-5.3.28.NC/build_unix
../dist/configure --enable-cxx --enable-java --prefix=/opt/db53
```

Build and Installation

You can now build and install Berkeley DB as follows:

Copyright 2016, ZeroC, Inc.

```
make
...
sudo make install
```

On a 64-bit platform, you need to create a `lib64` symbolic link to the `lib` directory:

```
cd /opt/db53
ln -s lib lib64
```

mcpp

Git Repository

Clone the `zeroc-inc/mcpp` repository hosted on `GitHub`:

```
git clone https://github.com/zeroc-ice/mcpp.git
```

Build and Installation

Follow the instructions in README.md.

Copyright 2016, ZeroC, Inc.

Building Berkeley DB on macOS

This page describes how to build Berkeley DB on macOS for use with Ice.

On this page:

- Berkeley DB
  - Source Archive
  - Patch
  - Configuration
  - Build and Installation

Berkeley DB

Source Archive

Download the Berkeley DB 5.3.28 source archive from Oracle or ZeroC. You may use the distribution with encryption or without encryption. Since Ice does not use Berkeley DB's encryption features, we recommend downloading the no-encryption (-NC) source archive:

```
wget https://zeroc.com/download/berkeley-db/db-5.3.28.NC.tar.gz
```

Then unpack this archive:

```
tar xzf db-5.3.28.NC.tar.gz
```

Patch

Download the following patch: https://zeroc.com/download/berkeley-db/berkeley-db.5.3.28.patch

Apply this patch from the main directory of the berkeley-db source distribution:

```
cd db-5.3.28.NC
patch -p0 < <path to patch>/berkeley-db.5.3.28.patch
```

Configuration

Berkeley DB must be configured with C++ support enabled. If you intend to use Ice for Java with Berkeley DB, you must also enable Java support. For example:

```
cd db-5.3.28.NC/build_unix
../dist/configure --enable-cxx --enable-java --prefix=/opt/db53
```

Build and Installation

You can now build and install Berkeley DB as follows:

```
make
...
sudo make install
```


Copyright 2016, ZeroC, Inc.

Ice-E Release Notes

Ice-E brings Ice to embedded devices, in particular Linux/ARM devices. Ice-E adds cross-compilation capabilities to Ice: with Ice plus Ice-E, you can develop C++ applications for Linux/ARM systems on your Linux x86_64 workstation.

The release notes provide information about an Ice-E release, including descriptions of significant new features and changes as well as instructions for upgrading from an earlier release.

Topics

- Supported Platforms for Ice-E 3.6.3
- New Features in Ice-E 3.6.3
- Ice-E Feature Set
- Ice-E Application Notes
- Upgrading your Ice-E Application
- Using the Ice-E Debian Binary Distribution
- Using Ice-E with Yocto

Copyright 2016, ZeroC, Inc.

Supported Platforms for Ice-E 3.6.3

Ice-E for C++

Toolchain	Development Platform	Run-Time Platform
Poky Yocto 2.1 (g++ 5.3)	Ubuntu 16.04 xenial amd64	Linux built using Yocto 2.1 on ARMv7-A
g++ 4.9	Debian 8 jessie amd64	Debian 8 jessie amd64
arm-linux-gnueabihf-g++-4.9	Debian 8 jessie amd64	Debian 8 jessie armhf

Ice-E for Python

Platform	Environment	Architecture
Linux built using Yocto 2.1	Python 2.7	ARMv7-A
Debian 8 jessie	Python 2.7	armhf

Copyright 2016, ZeroC, Inc.

New Features in Ice-E 3.6.3

This page outlines changes and improvements in this release that may affect the operation of your applications or have an impact on your source code.

Changes in Ice-E 3.6.3

- Upgraded to Ice 3.6.3.

Changes in Ice-E 3.6.2

- Added support for Freeze and IceStorm (with shared libraries only)
- Added support for Debian 8 amd64 as development and run-time environment
- Added support for Debian 8 armhf as target run-time environment
- Dropped support for Ubuntu 15.04 amd64 as development and runtime environment
- Dropped support for Debian 7.8 armhf as target run-time environment

Changes in Ice-E 3.6.1

- Added support for shared C++ libraries in addition to static C++ libraries
- Added support for IceBox for C++ (use with shared libraries).
- Added support for Yocto 1.7 and 1.8.
- The Glacier2 package for Debian 7.8 armhf was renamed `zeroc-glacier2` (it was named `zeroc-icee-glacier2` previously); the `glacier2router` is now linked with shared libraries provided by `libzeroc-ice3.6` for Debian 7.8 armhf.
- The statically-linked Glacier2 package for Ubuntu 15.04 `amd64` (`zeroc-icee-glacier2`) was removed. Use instead the `zeroc-glacier2` package for `amd64` provided by the Ice 3.6.1 distribution.
- The Python wheel for Debian 7.8 armhf was replaced by an apt package, `libzeroc-ice-python2.7`.

Changes in Ice-E 3.6.0

- Ice-E is now fully in-sync with Ice:
  - Ice-E uses the same developer kit as Ice: the Slice files, C++ header files, and Slice compiler (slice2cpp) are shared by Ice and Ice-E.
  - The Ice-E and Ice versions are now the same.
  - The preferred method to install Ice-E on Ubuntu is now through apt repositories.
- IceDiscovery, IceLocatorDiscovery, IceSSL, IceStringConverter and IceUtil are all included in the Ice static library.

Copyright 2016, ZeroC, Inc.

Ice-E Feature Set

Ice-E 3.6.3 is based on the Ice 3.6.3 release. See the Ice 3.6.3 Release Notes for more information on Ice.

Ice-E and Ice provide the same API and implement the same Ice protocol. Here is a brief comparison of these two products:

Feature	Ice-E	Ice
C++, Python	✓	✓
Java, JavaScript, Objective-C, PHP, Ruby	✗ (1)	✓
Ice transports: tcp, ssl, ws, wss, udp	✓	✓
Client libraries for Ice services (IceGrid, IceStorm etc.)	✓	✓
IceBox	✓ (2)	✓
Glacier2 firewall traversal service	✓ (2)	✓
Freeze persistent storage service	✓ (3)	✓
IceStorm publish-subscribe service	✓ (3)	✓
Other Ice services (IceGrid, IcePatch2)	✗	✓
Supports cross-compilation to Linux/ARM	✓	✗
Static libraries	✓	✗
Shared libraries	✓	✓

(1) Ice-E is supported only for C++ and Python. It is likely other Ice language mappings would work on Ice-E Linux/ARM targets, but were not tested.

(2) IceBox and Glacier2 are built only with shared libraries.

(3) Freeze and IceStorm are built only as shared libraries, and are available only on Debian.

Copyright 2016, ZeroC, Inc.

Ice-E Application Notes

On this page:

- Using Plug-ins with Static Libraries

Using Plug-ins with Static Libraries

The `IceDiscovery`, `IceLocatorDiscovery`, `IceSSL`, and `IceStringConverter` plug-ins are included with the `libIce.a` static library. To use these plug-ins, your application must ensure the plug-ins are linked in by calling the plug-in's register function and linking with the Ice-E static library.

The following functions can be used for registering the plug-ins:

- `Ice::registerIceSSL`
- `Ice::registerIceDiscovery`
- `Ice::registerIceLocatorDiscovery`
- `Ice::registerIceStringConverter`

Refer to the installing section of each plug-in in the Ice Plugins documentation for additional information on these methods.

If you are developing your own plug-in, you will also need to register your plugin factory function using `Ice::registerPluginFactory`. Refer to the Plug-in API for more information on this function.

Copyright 2016, ZeroC, Inc.

Upgrading your Ice-E Application

Ice-E 3.6.3 is fully source and binary compatible with prior versions of Ice-E 3.6. You do not need to recompile anything when upgrading your application from Ice-E 3.6.x.

Ice-E 3.6 is not compatible with Ice-E 1.x. As part of your upgrade from Ice-E 1.x to Ice-E 3.6, you need to update your source code and recompile all your Slice files.

Copyright 2016, ZeroC, Inc.

Using the Ice-E Debian Binary Distribution

```
deb http://emdebian.org/tools/debian/ jessie main
```

This page provides important information for users of the Ice-E Debian binary distribution.

On this page:

- Overview of the Ice-E Debian binary distribution
  - DEB packages
- Installing the Ice-E binary distribution
  - Installing Ice-E on Debian 8 amd64
  - Installing Ice-E on Debian 8 armhf
  - Setting up your cross development environment to use Ice-E
- Using the Ice-E sample programs

Overview of the Ice-E Debian binary distribution

The Ice-E binary distribution provides packages to develop applications for Debian 8 (jessie) armhf and amd64. The only supported development platform is Debian 8; applications targeting Debian 8 armhf can be cross-compiled using a Debian 8 amd64 system.

Ice-E has been extensively tested on the BeagleBone Black board (Rev C) with the default Debian 8 (jessie) distribution for this board.

DEB packages

ZeroC provides the following DEB packages for Debian 8 (jessie) amd64 and armhf:

Package	Co-installable	Description
zeroc-ice-all-dev	no	Meta package that installs the entire Ice dev-kit
zeroc-ice-all-runtime	no	Meta package that installs the entire Ice run-time, including [illegible] libraries
libzeroc-freeze3.6	yes	Freeze for C++ run-time (shared libraries)
libzeroc-ice-dev	yes	Ice for C++ dev-kit, including the dev-kit forservices such as Freeze, Glacier2 and IceStorm. This dev-kit consists of header files, static libraries and .so symbolic links.
libzeroc-ice-python2.7	no	Ice for Python run-time
libzeroc-ice3.6	yes	Ice for C++ run-time (shared libraries)
libzeroc-icestorm3.6	yes	IceStorm pub-sub service for IceBox C++
zeroc-glacier2	no	Glacier2 firewall-traversal service
zeroc-ice-compilers	no	Slice compilers: slice2cpp, slice2freeze, slice2py
zeroc-ice-slice	no	Slice files
zeroc-ice-utils	no	Ice utilities and admin tools
zeroc-icebox	no	IceBox server for C++

A co-installable package is a package for which you can install several architectures (amd64, armhf...) at the same time.

Installing the Ice-E binary distribution

Copyright 2016, ZeroC, Inc.

Installing Ice-E on Debian 8 amd64

Debian 8 amd64 can be used as development platform when targeting Debian 8 amd64 and Debian 8 armhf.

Follow the steps below to install Ice-E on Debian amd64.

1. Add the armhf architecture:

```
sudo dpkg --add-architecture armhf
```

2. Install ZeroC's GPG key (otherwise you will get warnings when installing ZeroC packages):

```
sudo apt-key adv --keyserver hkp://pgp.mit.edu:80 --recv-keys 5E6DA83306132997
```

3. Add the Ice repository for Debian 8 to your system:

```
sudo sh -c 'echo "deb [arch=amd64,armhf] http://zeroc.com/download/apt/debian8
stable main" > /etc/apt/sources.list.d/ice.list'
```

4. Update the package list:

```
sudo apt-get update
```

the `apt-get update` command will show some warnings for repositories that do not support the newly-added architecture. To avoid these warnings you can edit your repository lists and set the architectures supported by each repository (see "Setting up apt sources" in the Multiarch HOWTO).

5. Install the `amd64` packages:

```
sudo apt-get install zeroc-ice-all-runtime zeroc-ice-all-dev
```

6. Install the `armhf` cross development packages:

```
sudo apt-get install libzeroc-ice-dev:armhf
```

Refer to the package summary if you would like to install fewer packages.

Installing Ice-E on Debian 8 armhf

Debian 8 armhf can be used as development platform when targeting Debian 8 armhf. It is however more common to cross-compile on Debian 8 amd64 and target Debian 8 armhf.

Follow the steps below to install Ice-E on Debian armhf.

1. Install ZeroC's GPG key (otherwise you will get warnings when installing ZeroC packages):

Copyright 2016, ZeroC, Inc.

```
sudo apt-key adv --keyserver hkp://pgp.mit.edu:80 --recv-keys 5E6DA83306132997
```

2. Add the Ice repository for Debian 8 to your system:

```
sudo sh -c 'echo "deb http://zeroc.com/download/apt/debian8 stable main" >
/etc/apt/sources.list.d/ice.list'
```

3. Update the package list:

```
sudo apt-get update
```

4. Install the `armhf` run-time packages:

```
sudo apt-get install zeroc-ice-all-runtime
```

You can also include the full dev-kit on Debian armhf with:

```
sudo apt-get install zeroc-ice-all-dev
```

Refer to the package summary if you would like to install fewer packages.

Setting up your cross development environment to use Ice-E

This section shows how to setup your cross development environment on Debian 8 (Jessie) amd64 for developing applications targeting Debian 8 (Jessie) armhf. The required packages are not included in Jessie so you need to install them from the Debian Cross-toolchains repository.

1. Install Emdebian's key (otherwise you will get warnings when installing the packages):

```
sudo apt-key adv --keyserver pgpkeys.mit.edu --recv-keys 7DE089671804772E
```

2. Add the Debian Cross toolchains repository:

```
sudo sh -c 'echo "deb http://emdebian.org/tools/debian/ jessie main" >
/etc/apt/sources.list.d/crosstools.list'
```

If you have not already added the armhf architecture, do so now:

```
 sudo dpkg --add-architecture armhf
```

3. Update the package list:

Copyright 2016, ZeroC, Inc.

```
sudo apt-get update
```

4. Install cross compiler

```
sudo apt-get install g++-4.9-arm-linux-gnueabihf
```

When cross compiling Ice-E applications for Debian 8 armrhf you must use the `arm-linux-gnueabihf-g++-4.9` compiler instead of `g++`.

Sample compile command:

```
arm-linux-gnueabihf-g++-4.9 -pthread Hello.cpp
```

Ice-E applications can link with either the shared Ice libraries or with the static Ice library.

The static Ice library includes the symbols for Ice, IceUtil, IceSSL, IceDiscovery and IceLocatorDiscovery.

Use the following command to link your application with the shared Ice libraries:

```
arm-linux-gnueabihf-g++-4.9 -o client Hello.o Client.o -Wl,-Bdynamic -lIce
-lIceUtil -pthread
```

Use the following command to link your application with the static Ice library (`libIce.a`):

```
arm-linux-gnueabihf-g++-4.9 -o client Hello.o Client.o -Wl,-Bstatic -lIce
-Wl,-Bdynamic -lcrypto -lssl -pthread -lbz2 -ldl -lrt
```

For applications compiled with `--std=c++11`, add `++11` to all library names, for example `-lIce++11 -lIceUtil++11`.

Using the Ice-E sample programs

Sample programs are available in a separate GitHub repository. Simply clone this repository and use its 3.6 branch:

```
git clone -b 3.6 https://github.com/zeroc-ice/ice-demos.git
cd ice-demos
```

The Ice-E demos for C++ and Python are located in the `cpp` and `python` directories respectively. Review the build instructions for more information.

Copyright 2016, ZeroC, Inc.

Using Ice-E with Yocto

This page provides important information for users of the Ice-E binary distribution.

On this page:

- Overview of Ice-E with Yocto
  - Development Boards
  - Meta Layer
- Installing the Meta Layer
- Using the Meta Layer to install Ice-E
- Setting up your cross development environment to use Ice-E
- Using IceSSL with Platform CAs
- Using the Ice-E sample programs

Overview of Ice-E with Yocto

The Yocto Project allows you to create a custom Linux distribution for your embedded product. With ZeroC's meta layer for Yocto, you can easily include Ice for C++, Ice for Python and/or the Glacier2 service in this custom Linux distribution.

Development Boards

Ice-E was tested on:

Board	Yocto version
Freescale SABRE SD	2.1 (Krogoth) with the Freescale Community BSP
BeagleBone Black board (Rev C)	2.1 (Krogoth)

Meta Layer

Refer to the README.md file in the zeroc-ice/meta-zeroc repository.

Installing the Meta Layer

Installing the meta layer for ZeroC products is simply a matter of copying the contents of the repository onto your file system. For example, assuming you are using the poky repository, download the meta layer archive from the GitHub page and extract it alongside the other directories labeled `meta-<name>`. Once this is done, source the poky environment configuration script and add the `meta-zeroc` layer to the `BBLAYERS` variable in your `bblayers.conf` file:

```
BBLAYERS ?= " \
  /home/zeroc/dev/poky/meta \
  /home/zeroc/dev/poky/meta-yocto \
  /home/zeroc/dev/poky/meta-yocto-bsp \
  /home/zeroc/dev/poky/meta-zeroc \
  "
```

Using the Meta Layer to install Ice-E

To include Ice-E in your image, add the package from the `zeroc-ice` recipe you wish to use to your `local.conf` file:

Copyright 2016, ZeroC, Inc.

```
# Install Ice for C++ dynamic libraries and Ice for Python to image
IMAGE_INSTALL_append = " zeroc-ice zeroc-ice-python"

# Add the development package to the SDK
TOOLCHAIN_TARGET_TASK_append = " zeroc-ice-dev zeroc-ice-staticdev"

# Add the development package to the Native SDK
TOOLCHAIN_HOST_TASK_append  = " nativesdk-zeroc-ice-dev"
```

Setting up your cross development environment to use Ice-E

With the variables set in `local.conf` as shown above, you can generate an SDK for your image with the following command:

Create SDK for the core-image-minimal Image

```
bitbake core-image-minimal -c populate_sdk
```

Once complete, this SDK can be found in `tmp/deploy/sdk`. Execute the `.sh` file to install. By default your SDK will be installed into `/opt/poky/<version>/`. You can then source the cross-development environment as follows:

```
source /opt/poky/2.1/environment-setup-cortexa9hf-vfp-neon-poky-linux-gnueabi
```

Using IceSSL with Platform CAs

To use the `IceSSL.UsePlatformCAs` property you will need to install the `ca-certificates` package as part of your image. Since OpenSSL does not have its default `CAfile` (where IceSSL looks for the default platform CAs) set to the certificate bundle installed from `ca-certifcates` you will need to set an additional IceSSL property:

```
IceSSL.CAs=/etc/ssl/certs/ca-certificates.crt
```

Using the Ice-E sample programs

Sample programs are available in the zeroc-ice/ice-demos repository:

```
git clone -b 3.6 https://github.com/zeroc-ice/ice-demos.git
cd ice-demos
```

The Ice-E demos for C++ and Python are located in the `cpp` and `python` directories respectively. Review the build instructions for more information.

Copyright 2016, ZeroC, Inc.

Ice Touch Release Notes

The release notes provide information about an Ice Touch release, including descriptions of significant new features and changes as well as instructions for upgrading from an earlier release.

Topics

- Supported Platforms for Ice Touch 3.6.3
- New Features in Ice Touch 3.6.3
- Ice Touch Feature Set
- Upgrading your Ice Touch Application
- Using the Ice Touch Binary Distribution

Copyright 2016, ZeroC, Inc.

Supported Platforms for Ice Touch 3.6.3

Compiler	SDK	Development Platform	Run-Time Platform
Xcode 8.0	macOS 10.12 iOS 10.0	macOS 10.12	OS X 10.11 or greater iOS 8.0 or greater (armv7, armv7s, arm64), iOS Simulator 8.0 or greater
Xcode 7.3	OS X 10.11, macOS 10.12 iOS 9.0	OS X 10.11, macOS 10.12	OS X 10.11 or greater iOS 8.0 or greater (armv7, armv7s, arm64), iOS Simulator 8.0 or greater

Copyright 2016, ZeroC, Inc.

New Features in Ice Touch 3.6.3

This page outlines changes and improvements in this release that may affect the operation of your applications or have an impact on your source code.

Changes in Ice Touch 3.6.3

- Upgraded to Ice 3.6.3.

Changes in Ice Touch 3.6.2

- Upgraded to Ice 3.6.2.

Changes in Ice Touch 3.6.1

- Upgraded to Ice 3.6.1.
- Added compatibility for Xcode 7, OS X El Capitan (10.11), and iOS 9:
  - Ice Touch libraries built using Xcode 7 and the iOS 9.0 SDK incorporate LLVM bitcode. Starting with Xcode 7, bitcode is required for libraries used in iOS applications.
  - You need to build Ice Touch from sources with Xcode 7 as long as the Ice Touch binary distribution is built with Xcode 6.

Build Ice Touch using Xcode 7

```
# Select Xcode 7 beta
sudo xcode-select -s /Applications/Xcode-beta.app/Contents/Developer/

# Install Homebrew tap (only necessary if you have not done so before)
brew tap zeroc-ice/tap

# Install from source using Homebrew
brew update
brew install icetouch36 --build-from-source
```

Changes in Ice Touch 3.6.0

- Ice Touch is now fully in-sync with Ice:
  - The Ice Touch Objective-C mapping was moved to the ice repository.
  - The Ice Touch and Ice versions are now the same.
  - The preferred method to install Ice Touch is now through Homebrew.
- Ice Touch is installed under `/usr/local`; it is no longer installed in `/Library/Developer`.
- Previous versions of Ice Touch included a Xcode plug-in. This plug-in has been replaced by the Ice Builder for Xcode; this Ice Builder is a separate package independent of Ice Touch.

Copyright 2016, ZeroC, Inc.

Ice Touch Feature Set

On this page:

- Comparing Ice Touch and Ice
- Ice Touch language mappings
- Ice Touch limitations
- Support for Ice Touch background applications

Comparing Ice Touch and Ice

Ice Touch 3.6.3 is based on the Ice 3.6.3 release. See the Ice 3.6.3 Release Notes for more information on Ice.

Ice Touch and Ice provide the same API and implement the same Ice protocol. Here is a brief comparison of these two products:

Feature	Ice Touch	Ice
C++, Objective-C	✓	✓
Java, JavaScript, PHP, Python, Ruby	✗	✓
Ice transports: tcp, ssl, ws, wss, udp	✓	✓
Client libraries for IceGrid, IceStorm and Glacier2	✓	✓
Ice services (IceGrid, IceStorm etc.)	✗	✓
Develop on OS X	✓	✓
Develop with Xcode and Ice Builder	✓	✓
Target OS X	✓	✓
Target iOS	✓	✗
Static libraries	✓	✗
Shared libraries	✗	✓

Ice Touch language mappings

Ice Touch supports the following mappings:

- Objective-C mapping
- C++ mapping

Ice Touch limitations

Ice Touch currently lacks support for the `Ice::Application` and `Ice::Service` helper classes.

Ice Touch has limited support for:

- SSL

  Ice Touch for OS X uses the Ice for C++ SSL protocol plug-in based on OS X Secure Transport. For iOS devices, Ice Touch SSL provides only a subset of this functionality. Ice Touch servers on iOS cannot authenticate SSL clients for instance. The IceSSL property reference provides complete details.

- UDP
  On iOS, UDP requests do not transparently establish a cellular connection.

Copyright 2016, ZeroC, Inc.

Support for Ice Touch background applications

Ice Touch supports VoIP applications as documented in Executing Code in the Background.

Follow these steps to ensure your application is correctly configured:

- Add the `voip` flag to the `UIBackgroundModes` key of your application's `Info.plist`.
- Set the configuration property `Ice.Voip` to `1` in your communicator configuration properties. As described in the Configuring Sockets for VoIP Usage, this causes the Ice run time to set the `kCFStreamNetworkServiceType` property to `kCFStreamNetwo rkServiceTypeVoIP` for all sockets.
- Before moving your application to the background, call the `setKeepAliveTimeout:handler:` method to specify the frequency at which your application must be woken to maintain your service.

The sample application `objective-c/IceTouch/iPhone/voip` provides a demonstration of these steps.

Copyright 2016, ZeroC, Inc.

Upgrading your Ice Touch Application

Ice Touch backward compatibility

Ice Touch 3.6.3 is fully source and binary compatible with prior versions of Ice Touch 3.6: you do not need to recompile anything when upgrading from Ice Touch 3.6.x.

Ice Touch 3.6 is not compatible with Ice Touch 1.x. As part of your upgrade from Ice Touch 1.x to Ice Touch 3.6, you need to recompile your Slice files and rebuild your application.

Xcode project settings

For Xcode iOS and OS X applications, you need to update the project setting "Additional SDKs" to match the location of the new Ice Touch SDK installation. You should use:

- `/usr/local/lib/IceTouch/ObjC/$(PLATFORM_NAME).sdk` for the Objective-C SDK
- `/usr/local/lib/IceTouch/Cpp/$(PLATFORM_NAME).sdk` for the C++ SDK

Copyright 2016, ZeroC, Inc.

Using the Ice Touch Binary Distribution

This page provides important information for users of the Ice Touch binary distribution.

On this page:

- Overview of the Ice Touch binary distribution
- Installing the Ice Touch binary distribution
- Setting up your OS X environment to use Ice Touch
  - Using Xcode SDKs
- Using the Ice Touch sample programs

Overview of the Ice Touch binary distribution

Ice Touch is provided through Homebrew. The `icetouch` formula includes the following components:

- Xcode Objective-C SDKs for OS X, iPhone iOS, iPhone Simulator
- Xcode C++ SDKs for OS X, iPhone iOS, iPhone Simulator

More information about Homebrew is available at http://brew.sh.

Installing the Ice Touch binary distribution

Using Homebrew, you can install the distribution with this command:

```
brew tap zeroc-ice/tap
brew install icetouch36
```

The binaries in this distribution are 64-bit binaries for OS X and fat binaries with support for ARMv7, ARMv7s and ARM64 on iOS.

Setting up your OS X environment to use Ice Touch

Using Xcode SDKs

The Xcode SDKs requires the Ice Builder for Xcode package to be installed. Please refer to the Ice Builder for Xcode README.md for instructions on how to install it.

For Cocoa and iPhone applications, you must add the appropriate directory to the *Additional SDKs* setting in your Xcode project:

Objective-C SDK	`/usr/local/IceTouch/ObjC/$(PLATFORM_NAME).sdk`
C++ SDK	`/usr/local/IceTouch/Cpp/$(PLATFORM_NAME).sdk`

In addition, when creating a new iPhone Xcode project, you must set the Code Signing Resource Rules Path to:

`$(SDKROOT)/ResourceRules.plist`

You must also add the following to the `Frameworks` folder:

```
CFNetwork.framework
Security.framework
Foundation.framework
```

Using the Ice Touch sample programs

Sample programs are available in a separate GitHub repository. Simply clone this repository and use its 3.6 branch:

Copyright 2016, ZeroC, Inc.

```
git clone -b 3.6 https://github.com/zeroc-ice/ice-demos.git
cd ice-demos
```

The Ice Touch demos are located in the `cpp/IceTouch` and `objective-c/IceTouch` directories.

Copyright 2016, ZeroC, Inc.

Slice API Reference

Modules

Ice — The Ice core library.
Freeze — Freeze provides automatic persistence for Ice servants.
Glacier2 — Glacier2 is a firewall solution for Ice.
IceBox — IceBox is an application server specifically for Ice applications.
IceGrid — IceGrid is a server activation and deployment tool.
IceLocatorDiscovery — IceLocatorDiscovery is an Ice plug-in that enables the discovery of IceGrid and custom locators via UDP multicast.
IceMX — The Ice Management eXtension facility.
IcePatch2 — IcePatch can be used to update file hierarchies in a simple and efficient manner.
IceSSL — IceSSL provides a secure transport for Ice.
IceStorm — A messaging service with support for federation.

Copyright 2016, ZeroC, Inc.

Freeze Slice API

Freeze

Overview

module Freeze

Freeze provides automatic persistence for Ice servants.

Interface Index

BackgroundSaveEvictor — A background-save evictor is an evictor that saves updates asynchronously in a background thread.
Connection — A connection to a database (database environment with Berkeley DB).
Evictor — An automatic Ice object persistence manager, based on the evictor pattern.
EvictorIterator — An iterator for the objects managed by the evictor.
ServantInitializer — A servant initializer provides the application with an opportunity to perform custom servant initialization.
Transaction — A transaction.
TransactionalEvictor — A transactional evictor is an evictor that performs every single read-write operation within its own transaction.

Exception Index

DatabaseException — A Freeze database exception.
DeadlockException — A Freeze database deadlock exception.
EvictorDeactivatedException — This exception is raised if the evictor has been deactivated.
IndexNotFoundException — Exception raised when Freeze fails to locate an index.
InvalidPositionException — This Freeze Iterator is not on a valid position, for example this position has been erased.
NoSuchElementException — This exception is raised if there are no further elements in the iteration.
NotFoundException — A Freeze database exception, indicating that a database record could not be found.
TransactionAlreadyInProgressException — An attempt was made to start a transaction while a previous transaction has not yet been committed or rolled back.

Structure Index

CatalogData — The catalog keeps information about Freeze Maps and Freeze evictors in a Berkeley Db environment.
ObjectRecord — ObjectRecord is the value-type for the persistent maps maintained by evictors when using Ice encoding version is 1.0.
Statistics — Evictors maintain statistics about each object, when using Ice encoding version 1.0.

Sequence Index

Key — A database key, represented as a sequence of bytes.
Value — A database value, represented as a sequence of bytes.

Sequences

sequence<byte> Key

A database key, represented as a sequence of bytes.

sequence<byte> Value

A database value, represented as a sequence of bytes.

Copyright 2016, ZeroC, Inc.

Freeze-BackgroundSaveEvictor

Freeze::BackgroundSaveEvictor

Overview

local interface BackgroundSaveEvictor extends Freeze::Evictor

A background-save evictor is an evictor that saves updates asynchronously in a background thread.

Operation Index

keep — Lock this object in the evictor cache.
keepFacet — Like keep, but with a facet.
release — Release a lock acquired by keep.
releaseFacet — Like release, but with a facet.

Operations

void keep(Ice::Identity id)

Lock this object in the evictor cache. This lock can be released by release or `remove`. release releases only one lock, while `remove` release s all the locks.

Parameters

`id` — The identity of the Ice object.

Exceptions

Ice::NotRegisteredException — Raised if this identity was not registered with the evictor.
Freeze::DatabaseException — Raised if a database failure occurred.

See Also

- keepFacet
- release
- `remove`

void keepFacet(Ice::Identity id, string facet)

Like keep, but with a facet. Calling `keep(id)` is equivalent to calling keepFacet with an empty facet.

Parameters

`id` — The identity of the Ice object.
`facet` — The facet. An empty facet means the default facet.

Exceptions

Ice::NotRegisteredException — Raised if this identity was not registered with the evictor.
Freeze::DatabaseException — Raised if a database failure occurred.

See Also

- keep
- releaseFacet
- `removeFacet`

void release(Ice::Identity id)

Release a lock acquired by keep. Once all the locks on an object have been released, the object is again subject to the normal eviction

Copyright 2016, ZeroC, Inc.

strategy.

Parameters

`id` — The identity of the Ice object.

Exceptions

Ice::NotRegisteredException — Raised if this object was not locked with keep or keepFacet.

See Also

- keepFacet
- release

void releaseFacet(**Ice::Identity** id, string facet)

Like release, but with a facet. Calling `release(id)` is equivalent to calling releaseFacet with an empty facet.

Parameters

`id` — The identity of the Ice object.
`facet` — The facet. An empty facet means the default facet.

Exceptions

Ice::NotRegisteredException — Raised if this object was not locked with keep or keepFacet.

See Also

- keep
- releaseFacet

Copyright 2016, ZeroC, Inc.

Freeze-CatalogData

Freeze::CatalogData

Overview

struct CatalogData

The catalog keeps information about Freeze Maps and Freeze evictors in a Berkeley Db environment. It is used by FreezeScript.

Data Member Index

evictor — True if this entry describes an evictor database, false if it describes a map database.
key — The Slice type for the database key.
value — The Slice type for the database value.

Data Members

bool evictor;

True if this entry describes an evictor database, false if it describes a map database.

string key;

The Slice type for the database key.

string value;

The Slice type for the database value.

Copyright 2016, ZeroC, Inc.

Freeze-Connection

Freeze::Connection

Overview

local interface Connection

A connection to a database (database environment with Berkeley DB). If you want to use a connection concurrently in multiple threads, you need to serialize access to this connection.

Used By

- Freeze::Transaction::getConnection

Operation Index

beginTransaction — Create a new transaction.
currentTransaction — Returns the transaction associated with this connection.
removeMapIndex — Remove an old unused Freeze Map index @throws IndexNotFoundException Raised if this index does not exist
close — Closes this connection.
getCommunicator — Returns the communicator associated with this connection.
getEncoding — Returns the encoding version used to encode the data.
getName — The name of the connected system (for example, the Berkeley DB environment).

Operations

Freeze::Transaction beginTransaction()

Create a new transaction. Only one transaction at a time can be associated with a connection.

Return Value

The new transaction.

Exceptions

Freeze::TransactionAlreadyInProgressException — Raised if a transaction is already associated with this connection.

Freeze::Transaction currentTransaction()

Returns the transaction associated with this connection.

Return Value

The current transaction if there is one, null otherwise.

void removeMapIndex(string mapName, string indexName)

Remove an old unused Freeze Map index

Exceptions

Freeze::IndexNotFoundException — Raised if this index does not exist

void close()

Closes this connection. If there is an associated transaction, it is rolled back.

Ice::Communicator getCommunicator()

Returns the communicator associated with this connection.

Copyright 2016, ZeroC, Inc.

Ice::EncodingVersion getEncoding()

Returns the encoding version used to encode the data.

string getName()

The name of the connected system (for example, the Berkeley DB environment).

Copyright 2016, ZeroC, Inc.

Freeze-DatabaseException

Freeze::DatabaseException

Overview

local exception DatabaseException

A Freeze database exception.

Derived Exceptions

- Freeze::DeadlockException
- Freeze::NotFoundException

See Also

- Freeze::Evictor
- Freeze::Connection

Data Member Index

message — A message describing the reason for the exception.

Data Members

string message;

A message describing the reason for the exception.

Copyright 2016, ZeroC, Inc.

Freeze-DeadlockException

Freeze::DeadlockException

Overview

local exception DeadlockException extends Freeze::DatabaseException

A Freeze database deadlock exception. Applications can react to this exception by aborting and trying the transaction again.

Data Member Index

tx — The transaction in which the deadlock occurred.

Data Members

Freeze::Transaction tx;

The transaction in which the deadlock occurred.

Copyright 2016, ZeroC, Inc.

Freeze-Evictor

Freeze::Evictor

Overview

local interface Evictor extends Ice::ServantLocator

An automatic Ice object persistence manager, based on the evictor pattern. The evictor is a servant locator implementation that stores the persistent state of its objects in a database. Any number of objects can be registered with an evictor, but only a configurable number of servants are active at a time. These active servants reside in a queue; the least recently used servant in the queue is the first to be evicted when a new servant is activated.

Derived Classes and Interfaces

- Freeze::BackgroundSaveEvictor
- Freeze::TransactionalEvictor

See Also

- Freeze::ServantInitializer

Operation Index

setSize — Set the size of the evictor's servant queue.
getSize — Get the size of the evictor's servant queue.
add — Add a servant to this evictor.
addFacet — Like add, but with a facet.
remove — Permanently destroy an Ice object.
removeFacet — Like remove, but with a facet.
hasObject — Returns true if the given identity is managed by the evictor with the default facet.
hasFacet — Like hasObject, but with a facet.
getIterator — Get an iterator for the identities managed by the evictor.

Operations

void setSize(int sz)

Set the size of the evictor's servant queue. This is the maximum number of servants the evictor keeps active. Requests to set the queue size to a value smaller than zero are ignored.

Parameters

`sz` — The size of the servant queue. If the evictor currently holds more than `sz` servants in its queue, it evicts enough servants to match the new size. Note that this operation can block if the new queue size is smaller than the current number of servants that are servicing requests. In this case, the operation waits until a sufficient number of servants complete their requests.

Exceptions

Freeze::EvictorDeactivatedException — Raised if a the evictor has been deactivated.

See Also

- getSize

int getSize()

Get the size of the evictor's servant queue.

Return Value

The size of the servant queue.

Copyright 2016, ZeroC, Inc.

Exceptions

Freeze::EvictorDeactivatedException — Raised if a the evictor has been deactivated.

See Also

- setSize

Object* add(Object servant, Ice::Identity id)

Add a servant to this evictor. The state of the servant passed to this operation will be saved in the evictor's persistent store.

Parameters

`servant` — The servant to add.
`id` — The identity of the Ice object that is implemented by the servant.

Return Value

A proxy that matches the given identity and this evictor's object adapter.

Exceptions

Ice::AlreadyRegisteredException — Raised if the evictor already has an object with this identity.
Freeze::DatabaseException — Raised if a database failure occurred.
Freeze::EvictorDeactivatedException — Raised if the evictor has been deactivated.

See Also

- addFacet
- remove
- removeFacet

Object* addFacet(Object servant, Ice::Identity id, string facet)

Like add, but with a facet. Calling `add(servant, id)` is equivalent to calling addFacet with an empty facet.

Parameters

`servant` — The servant to add.
`id` — The identity of the Ice object that is implemented by the servant.
`facet` — The facet. An empty facet means the default facet.

Return Value

A proxy that matches the given identity and this evictor's object adapter.

Exceptions

Ice::AlreadyRegisteredException — Raised if the evictor already has an object with this identity.
Freeze::DatabaseException — Raised if a database failure occurred.
Freeze::EvictorDeactivatedException — Raised if the evictor has been deactivated.

See Also

- add
- remove
- removeFacet

Object remove(Ice::Identity id)

Permanently destroy an Ice object.

Parameters

`id` — The identity of the Ice object.

Copyright 2016, ZeroC, Inc.

Return Value

The removed servant.

Exceptions

Ice::NotRegisteredException — Raised if this identity was not registered with the evictor.
Freeze::DatabaseException — Raised if a database failure occurred.
Freeze::EvictorDeactivatedException — Raised if the evictor has been deactivated.

See Also

- add
- removeFacet

Object removeFacet(Ice::Identity id, string facet)

Like remove, but with a facet. Calling `remove(id)` is equivalent to calling removeFacet with an empty facet.

Parameters

`id` — The identity of the Ice object.
`facet` — The facet. An empty facet means the default facet.

Return Value

The removed servant.

Exceptions

Ice::NotRegisteredException — Raised if this identity was not registered with the evictor.
Freeze::DatabaseException — Raised if a database failure occurred.
Freeze::EvictorDeactivatedException — Raised if the evictor has been deactivated.

See Also

- remove
- addFacet

bool hasObject(Ice::Identity id)

Returns true if the given identity is managed by the evictor with the default facet.

Return Value

true if the identity is managed by the evictor, false otherwise.

Exceptions

Freeze::DatabaseException — Raised if a database failure occurred.
Freeze::EvictorDeactivatedException — Raised if a the evictor has been deactivated.

bool hasFacet(Ice::Identity id, string facet)

Like hasObject, but with a facet. Calling `hasObject(id)` is equivalent to calling hasFacet with an empty facet.

Return Value

true if the identity is managed by the evictor for the given facet, false otherwise.

Exceptions

Freeze::DatabaseException — Raised if a database failure occurred.
Freeze::EvictorDeactivatedException — Raised if a the evictor has been deactivated.

Copyright 2016, ZeroC, Inc.

Freeze::EvictorIterator getIterator(string facet, int batchSize)

Get an iterator for the identities managed by the evictor.

Parameters

`facet` — The facet. An empty facet means the default facet.
`batchSize` — Internally, the Iterator retrieves the identities in batches of size batchSize. Selecting a small batchSize can have an adverse effect on performance.

Return Value

A new iterator.

Exceptions

Freeze::EvictorDeactivatedException — Raised if a the evictor has been deactivated.

Copyright 2016, ZeroC, Inc.

Freeze-EvictorDeactivatedException

Freeze::EvictorDeactivatedException

Overview

local exception EvictorDeactivatedException

This exception is raised if the evictor has been deactivated.

Copyright 2016, ZeroC, Inc.

Freeze-EvictorIterator

Freeze::EvictorIterator

Overview

local interface EvictorIterator

An iterator for the objects managed by the evictor. Note that an EvictorIterator is not thread-safe: the application needs to serialize access to a given EvictorIterator, for example by using it in just one thread.

Used By

- Freeze::Evictor::getIterator

See Also

- Freeze::Evictor

Operation Index

hasNext — Determines if the iteration has more elements.
next — Obtains the next identity in the iteration.

Operations

bool hasNext()

Determines if the iteration has more elements.

Return Value

True if the iterator has more elements, false otherwise.

Exceptions

Freeze::DatabaseException — Raised if a database failure occurs while retrieving a batch of objects.

Ice::Identity next()

Obtains the next identity in the iteration.

Return Value

The next identity in the iteration.

Exceptions

Freeze::NoSuchElementException — Raised if there is no further elements in the iteration.
Freeze::DatabaseException — Raised if a database failure occurs while retrieving a batch of objects.

Copyright 2016, ZeroC, Inc.

Freeze-IndexNotFoundException

Freeze::IndexNotFoundException

Overview

local exception IndexNotFoundException

Exception raised when Freeze fails to locate an index.

Data Member Index

mapName — The name of the map in which the index could not be found.
indexName — The name of the index.

Data Members

string mapName;

The name of the map in which the index could not be found.

string indexName;

The name of the index.

Copyright 2016, ZeroC, Inc.

Freeze-InvalidPositionException

Freeze::InvalidPositionException

Overview

local exception InvalidPositionException

This Freeze Iterator is not on a valid position, for example this position has been erased.

Copyright 2016, ZeroC, Inc.

Freeze-NoSuchElementException

Freeze::NoSuchElementException

Overview

local exception NoSuchElementException

This exception is raised if there are no further elements in the iteration.

Copyright 2016, ZeroC, Inc.

Freeze-NotFoundException

Freeze::NotFoundException

Overview

local exception NotFoundException extends Freeze::DatabaseException

A Freeze database exception, indicating that a database record could not be found.

Copyright 2016, ZeroC, Inc.

Freeze-ObjectRecord

Freeze::ObjectRecord

Overview

struct ObjectRecord

ObjectRecord is the value-type for the persistent maps maintained by evictors when using Ice encoding version is 1.0.

Data Member Index

servant — The servant implementing the object.
stats — The statistics for the object.

Data Members

Object servant;

The servant implementing the object.

Freeze::Statistics stats;

The statistics for the object.

Copyright 2016, ZeroC, Inc.

Freeze-ServantInitializer

Freeze::ServantInitializer

Overview

local interface ServantInitializer

A servant initializer provides the application with an opportunity to perform custom servant initialization.

See Also

- Freeze::Evictor

Operation Index

initialize — Called whenever the evictor creates a new servant.

Operations

void initialize(**Ice::ObjectAdapter** adapter, **Ice::Identity** identity, string facet, Object servant)

Called whenever the evictor creates a new servant. This operation allows application code to perform custom servant initialization after the servant has been created by the evictor and its persistent state has been restored.

Parameters

`adapter` — The object adapter in which the evictor is installed.
`identity` — The identity of the Ice object for which the servant was created.
`facet` — The facet. An empty facet means the default facet.
`servant` — The servant to initialize.

See Also

- Ice::Identity

Copyright 2016, ZeroC, Inc.

Freeze-Statistics

Freeze::Statistics

Overview

struct Statistics

Evictors maintain statistics about each object, when using Ice encoding version 1.0.

Used By

- Freeze::ObjectRecord::stats

Data Member Index

creationTime — The time the object was created, in milliseconds since Jan 1, 1970 0:00.
lastSaveTime — The time the object was last saved, in milliseconds relative to `creationTime`.
avgSaveTime — The average time between saves, in milliseconds.

Data Members

long creationTime;

The time the object was created, in milliseconds since Jan 1, 1970 0:00.

long lastSaveTime;

The time the object was last saved, in milliseconds relative to `creationTime`.

long avgSaveTime;

The average time between saves, in milliseconds.

Copyright 2016, ZeroC, Inc.

Freeze-Transaction

Freeze::Transaction

Overview

local interface Transaction

A transaction. If you want to use a transaction concurrently in multiple threads, you need to serialize access to this transaction.

Used By

- Freeze::Connection::beginTransaction
- Freeze::Connection::currentTransaction
- Freeze::DeadlockException::tx
- Freeze::TransactionalEvictor::getCurrentTransaction
- Freeze::TransactionalEvictor::setCurrentTransaction

Operation Index

commit — Commit this transaction.
rollback — Roll back this transaction.
getConnection — Get the connection associated with this Transaction

Operations

void commit()

Commit this transaction.

Exceptions

Freeze::DatabaseException — Raised if a database failure occurred.

void rollback()

Roll back this transaction.

Exceptions

Freeze::DatabaseException — Raised if a database failure occurred.

Freeze::Connection getConnection()

Get the connection associated with this Transaction

Copyright 2016, ZeroC, Inc.

Freeze-TransactionalEvictor

Freeze::TransactionalEvictor

Overview

local interface TransactionalEvictor extends Freeze::Evictor

A transactional evictor is an evictor that performs every single read-write operation within its own transaction.

Operation Index

getCurrentTransaction — Get the transaction associated with the calling thread.
setCurrentTransaction — Associate a transaction with the calling thread.

Operations

Freeze::Transaction getCurrentTransaction()

Get the transaction associated with the calling thread.

Return Value

The transaction associated with the calling thread.

void setCurrentTransaction(Freeze::Transaction tx)

Associate a transaction with the calling thread.

Parameters

`tx` — The transaction to associate with the calling thread.

Copyright 2016, ZeroC, Inc.

Freeze-TransactionAlreadyInProgressException

Freeze::TransactionAlreadyInProgressException

Overview

local exception TransactionAlreadyInProgressException

An attempt was made to start a transaction while a previous transaction has not yet been committed or rolled back.

Copyright 2016, ZeroC, Inc.

Glacier2 Slice API

Glacier2

Overview

["objc:prefix:GLACIER2"] module Glacier2

Glacier2 is a firewall solution for Ice. Glacier2 authenticates and filters client requests and allows callbacks to the client in a secure fashion. In combination with IceSSL, Glacier2 provides a security solution that is both non-intrusive and easy to configure.

Interface Index

IdentitySet — An object for managing the set of object identity constraints on a Session.
PermissionsVerifier — The Glacier2 permissions verifier.
Router — The Glacier2 specialization of the Ice::Router interface.
SSLPermissionsVerifier — The SSL Glacier2 permissions verifier.
SSLSessionManager — The session manager for SSL authenticated users that is responsible for managing Session objects.
Session — A client-visible session object, which is tied to the lifecycle of a Router.
SessionControl — An administrative session control object, which is tied to the lifecycle of a Session.
SessionManager — The session manager for username/password authenticated users that is responsible for managing Session objects.
StringSet — An object for managing the set of identity constraints for specific parts of object identity on a Session.

Exception Index

CannotCreateSessionException — This exception is raised if an attempt to create a new session failed.
PermissionDeniedException — This exception is raised if a client is denied the ability to create a session with the router.
SessionNotExistException — This exception is raised if a client tries to destroy a session with a router, but no session exists for the client.

Structure Index

SSLInfo — Information taken from an SSL connection used for permissions verification.

Copyright 2016, ZeroC, Inc.

Glacier2-CannotCreateSessionException

Glacier2::CannotCreateSessionException

Overview

["preserve-slice"] exception CannotCreateSessionException

This exception is raised if an attempt to create a new session failed.

Used By

- Glacier2::Router::createSessionFromSecureConnection
- Glacier2::Router::createSession
- Glacier2::SSLSessionManager::create
- Glacier2::SessionManager::create

Data Member Index

reason — The reason why session creation has failed.

Data Members

string reason;

The reason why session creation has failed.

Copyright 2016, ZeroC, Inc.

Glacier2-IdentitySet

Glacier2::IdentitySet

Overview

interface IdentitySet

An object for managing the set of object identity constraints on a Glacier2::Session.

See Also

- Glacier2::Session
- Glacier2::SessionControl

Operation Index

add — Add a sequence of Ice identities to this set of constraints.
remove — Remove a sequence of identities from this set of constraints.
get — Returns a sequence of identities describing the constraints in this set.

Operations

void add(**Ice::IdentitySeq** additions)

Add a sequence of Ice identities to this set of constraints. Order is not preserved and duplicates are implicitly removed.

Parameters

`additions` — The sequence of Ice identities to be added.

void remove(**Ice::IdentitySeq** deletions)

Remove a sequence of identities from this set of constraints. No errors are returned if an entry is not found.

Parameters

`deletions` — The sequence of Ice identities to be removed.

Ice::IdentitySeq get()

Returns a sequence of identities describing the constraints in this set.

Return Value

The sequence of Ice identities for this set.

Copyright 2016, ZeroC, Inc.

Glacier2-PermissionDeniedException

Glacier2::PermissionDeniedException

Overview

["preserve-slice"] exception PermissionDeniedException

This exception is raised if a client is denied the ability to create a session with the router.

Used By

- Glacier2::PermissionsVerifier::checkPermissions
- Glacier2::Router::createSessionFromSecureConnection
- Glacier2::Router::createSession
- Glacier2::SSLPermissionsVerifier::authorize

Data Member Index

reason — The reason why permission was denied.

Data Members

string reason;

The reason why permission was denied.

Copyright 2016, ZeroC, Inc.

Glacier2-PermissionsVerifier

Glacier2::PermissionsVerifier

Overview

interface PermissionsVerifier

The Glacier2 permissions verifier. This is called through the process of establishing a session.

See Also

- Glacier2::Router

Operation Index

checkPermissions — Check whether a user has permission to access the router.

Operations

["nonmutating", "format:sliced"] bool checkPermissions(string userId, string password, out string reason) throws Glacier2::PermissionDeniedException

Check whether a user has permission to access the router.

Parameters

`userId` — The user id for which to check permission.
`password` — The user's password.
`reason` — The reason why access was denied.

Return Value

True if access is granted, or false otherwise.

Exceptions

Glacier2::PermissionDeniedException — Raised if the user access is denied. This can be raised in place of returning false with a reason set in the reason out parameter.

Copyright 2016, ZeroC, Inc.

Glacier2-Router

Glacier2::Router

Overview

interface Router extends Ice::Router

The Glacier2 specialization of the Ice::Router interface.

Operation Index

getCategoryForClient — This category must be used in the identities of all of the client's callback objects.
createSession — Create a per-client session with the router.
createSessionFromSecureConnection — Create a per-client session with the router.
refreshSession — Keep the calling client's session with this router alive.
destroySession — Destroy the calling client's session with this router.
getSessionTimeout — Get the value of the session timeout.
getACMTimeout — Get the value of the ACM timeout.

Operations

["nonmutating"] string getCategoryForClient()

This category must be used in the identities of all of the client's callback objects. This is necessary in order for the router to forward callback requests to the intended client. If the Glacier2 server endpoints are not set, the returned category is an empty string.

Return Value

The category.

["amd", "format:sliced"] Glacier2::Session* createSession(string userId, string password) throws Glacier2::PermissionDeniedException, Glacier2::CannotCreateSessionException

Create a per-client session with the router. If a Glacier2::SessionManager has been installed, a proxy to a Glacier2::Session object is returned to the client. Otherwise, null is returned and only an internal session (i.e., not visible to the client) is created. If a session proxy is returned, it must be configured to route through the router that created it. This will happen automatically if the router is configured as the client's default router at the time the session proxy is created in the client process, otherwise the client must configure the session proxy explicitly.

Parameters

`userId` — The user id for which to check the password.
`password` — The password for the given user id.

Return Value

A proxy for the newly created session, or null if no Glacier2::SessionManager has been installed.

Exceptions

Glacier2::PermissionDeniedException — Raised if the password for the given user id is not correct, or if the user is not allowed access.
Glacier2::CannotCreateSessionException — Raised if the session cannot be created.

See Also

- Glacier2::Session
- Glacier2::SessionManager
- Glacier2::PermissionsVerifier

["amd", "format:sliced"] Glacier2::Session* createSessionFromSecureConnection() throws Glacier2::PermissionDeniedException, Glacier2::CannotCreateSessionException

Copyright 2016, ZeroC, Inc.

Create a per-client session with the router. The user is authenticated through the SSL certificates that have been associated with the connection. If a Glacier2::SessionManager has been installed, a proxy to a Glacier2::Session object is returned to the client. Otherwise, null is returned and only an internal session (i.e., not visible to the client) is created. If a session proxy is returned, it must be configured to route through the router that created it. This will happen automatically if the router is configured as the client's default router at the time the session proxy is created in the client process, otherwise the client must configure the session proxy explicitly.

Return Value

A proxy for the newly created session, or null if no Glacier2::SessionManager has been installed.

Exceptions

Glacier2::PermissionDeniedException — Raised if the user cannot be authenticated or if the user is not allowed access.
Glacier2::CannotCreateSessionException — Raised if the session cannot be created.

See Also

- Glacier2::Session
- Glacier2::SessionManager
- Glacier2::PermissionsVerifier

["amd"] void refreshSession() throws Glacier2::SessionNotExistException

Keep the calling client's session with this router alive.

Exceptions

Glacier2::SessionNotExistException — Raised if no session exists for the calling client.

void destroySession() throws Glacier2::SessionNotExistException

Destroy the calling client's session with this router.

Exceptions

Glacier2::SessionNotExistException — Raised if no session exists for the calling client.

["nonmutating"] long getSessionTimeout()

Get the value of the session timeout. Sessions are destroyed if they see no activity for this period of time.

Return Value

The timeout (in seconds).

["nonmutating"] int getACMTimeout()

Get the value of the ACM timeout. Clients supporting connection heartbeats can enable them instead of explicitly sending keep alives requests. NOTE: This method is only available since Ice 3.6.

Return Value

The timeout (in seconds).

Copyright 2016, ZeroC, Inc.

Glacier2-Session

Glacier2::Session

Overview

interface Session

A client-visible session object, which is tied to the lifecycle of a Glacier2::Router.

Derived Classes and Interfaces

- IceGrid::AdminSession
- IceGrid::Session

See Also

- Glacier2::Router
- Glacier2::SessionManager

Operation Index

destroy — Destroy the session.

Operations

void destroy()

Destroy the session. This is called automatically when the router is destroyed.

Copyright 2016, ZeroC, Inc.

Glacier2-SessionControl

Glacier2::SessionControl

Overview

interface SessionControl

An administrative session control object, which is tied to the lifecycle of a Glacier2::Session.

See Also

- Glacier2::Session

Operation Index

categories — Access the object that manages the allowable categories for object identities for this session.
adapterIds — Access the object that manages the allowable adapter identities for objects for this session.
identities — Access the object that manages the allowable object identities for this session.
getSessionTimeout — Get the session timeout.
destroy — Destroy the associated session.

Operations

Glacier2::StringSet* categories()

Access the object that manages the allowable categories for object identities for this session.

Return Value

A StringSet object.

Glacier2::StringSet* adapterIds()

Access the object that manages the allowable adapter identities for objects for this session.

Return Value

A StringSet object.

Glacier2::IdentitySet* identities()

Access the object that manages the allowable object identities for this session.

Return Value

An IdentitySet object.

int getSessionTimeout()

Get the session timeout.

Return Value

The timeout.

void destroy()

Destroy the associated session.

Copyright 2016, ZeroC, Inc.

Glacier2-SessionManager

Glacier2::SessionManager

Overview

interface SessionManager

The session manager for username/password authenticated users that is responsible for managing Glacier2::Session objects. New session objects are created by the Glacier2::Router object calling on an application-provided session manager. If no session manager is provided by the application, no client-visible sessions are passed to the client.

See Also

- Glacier2::Router
- Glacier2::Session

Operation Index

create — Create a new session.

Operations

["format:sliced"] Glacier2::Session* create(string userId, Glacier2::SessionControl* control) throws Glacier2::CannotCreateSessionException

Create a new session.

Parameters

`userId` — The user id for the session.
`control` — A proxy to the session control object.

Return Value

A proxy to the newly created session.

Exceptions

Glacier2::CannotCreateSessionException — Raised if the session cannot be created.

Copyright 2016, ZeroC, Inc.

Glacier2-SessionNotExistException

Glacier2::SessionNotExistException

Overview

exception SessionNotExistException

This exception is raised if a client tries to destroy a session with a router, but no session exists for the client.

Used By

- Glacier2::Router::destroySession
- Glacier2::Router::refreshSession

See Also

- Glacier2::Router::destroySession

Copyright 2016, ZeroC, Inc.

Glacier2-SSLInfo

Glacier2::SSLInfo

Overview

struct SSLInfo

Information taken from an SSL connection used for permissions verification.

Used By

- Glacier2::SSLPermissionsVerifier::authorize
- Glacier2::SSLSessionManager::create

See Also

- Glacier2::PermissionsVerifier

Data Member Index

remoteHost — The remote host.
remotePort — The remote port.
localHost — The router's host.
localPort — The router's port.
cipher — The negotiated cipher suite.
certs — The certificate chain.

Data Members

string remoteHost;

The remote host.

int remotePort;

The remote port.

string localHost;

The router's host.

int localPort;

The router's port.

string cipher;

The negotiated cipher suite.

Ice::StringSeq certs;

The certificate chain.

Copyright 2016, ZeroC, Inc.

Glacier2-SSLPermissionsVerifier

Glacier2::SSLPermissionsVerifier

Overview

interface SSLPermissionsVerifier

The SSL Glacier2 permissions verifier. This is called through the process of establishing a session.

See Also

- Glacier2::Router

Operation Index

authorize — Check whether a user has permission to access the router.

Operations

["nonmutating", "format:sliced"] bool authorize(**Glacier2::SSLInfo** info, out string reason) throws **Glacier2::PermissionDeniedException**

Check whether a user has permission to access the router.

Parameters

`info` — The SSL information.
`reason` — The reason why access was denied.

Return Value

True if access is granted, or false otherwise.

Exceptions

Glacier2::PermissionDeniedException — Raised if the user access is denied. This can be raised in place of returning false with a reason set in the reason out parameter.

See Also

- Glacier2::SSLInfo

Copyright 2016, ZeroC, Inc.

Glacier2-SSLSessionManager

Glacier2::SSLSessionManager

Overview

interface SSLSessionManager

The session manager for SSL authenticated users that is responsible for managing Glacier2::Session objects. New session objects are created by the Glacier2::Router object calling on an application-provided session manager. If no session manager is provided by the application, no client-visible sessions are passed to the client.

See Also

- Glacier2::Router
- Glacier2::Session

Operation Index

create — Create a new session.

Operations

["format:sliced"] Glacier2::Session* create(Glacier2::SSLInfo info, Glacier2::SessionControl* control) throws Glacier2::CannotCreateSessionException

Create a new session.

Parameters

`info` — The SSL info.
`control` — A proxy to the session control object.

Return Value

A proxy to the newly created session.

Exceptions

Glacier2::CannotCreateSessionException — Raised if the session cannot be created.

Copyright 2016, ZeroC, Inc.

Glacier2-StringSet

Glacier2::StringSet

Overview

interface StringSet

An object for managing the set of identity constraints for specific parts of object identity on a Glacier2::Session.

See Also

- Glacier2::Session
- Glacier2::SessionControl

Operation Index

add — Add a sequence of strings to this set of constraints.
remove — Remove a sequence of strings from this set of constraints.
get — Returns a sequence of strings describing the constraints in this set.

Operations

void add(**Ice::StringSeq** additions)

Add a sequence of strings to this set of constraints. Order is not preserved and duplicates are implicitly removed.

Parameters

`additions` — The sequence of strings to be added.

void remove(**Ice::StringSeq** deletions)

Remove a sequence of strings from this set of constraints. No errors are returned if an entry is not found.

Parameters

`deletions` — The sequence of strings to be removed.

Ice::StringSeq get()

Returns a sequence of strings describing the constraints in this set.

Return Value

The sequence of strings for this set.

Copyright 2016, ZeroC, Inc.

IceBox Slice API

IceBox

Overview

module IceBox

IceBox is an application server specifically for Ice applications. IceBox can easily run and administer Ice services that are dynamically loaded as a DLL, shared library, or Java class.

Interface Index

Service — An application service managed by a ServiceManager.
ServiceManager — Administers a set of Service instances.
ServiceObserver — An Observer interface implemented by admin clients interested in the status of services @see ServiceManager

Exception Index

AlreadyStartedException — This exception is thrown if an attempt is made to start an already-started service.
AlreadyStoppedException — This exception is thrown if an attempt is made to stop an already-stopped service.
FailureException — This exception is a general failure notification.
NoSuchServiceException — This exception is thrown if a service name does not refer to an existing service.

Copyright 2016, ZeroC, Inc.

IceBox-AlreadyStartedException

IceBox::AlreadyStartedException

Overview

exception AlreadyStartedException

This exception is thrown if an attempt is made to start an already-started service.

Used By

- IceBox::ServiceManager::startService

Copyright 2016, ZeroC, Inc.

IceBox-AlreadyStoppedException

IceBox::AlreadyStoppedException

Overview

exception AlreadyStoppedException

This exception is thrown if an attempt is made to stop an already-stopped service.

Used By

- IceBox::ServiceManager::stopService

Copyright 2016, ZeroC, Inc.

IceBox-FailureException

IceBox::FailureException

Overview

local exception FailureException

This exception is a general failure notification. It is thrown for errors such as a service encountering an error during initialization, or the service manager being unable to load a service executable.

Data Member Index

reason — The reason for the failure.

Data Members

string reason;

The reason for the failure.

Copyright 2016, ZeroC, Inc.

IceBox-NoSuchServiceException

IceBox::NoSuchServiceException

Overview

exception NoSuchServiceException

This exception is thrown if a service name does not refer to an existing service.

Used By

- IceBox::ServiceManager::startService
- IceBox::ServiceManager::stopService

Copyright 2016, ZeroC, Inc.

IceBox-Service

IceBox::Service

Overview

local interface Service

An application service managed by a IceBox::ServiceManager.

Operation Index

start — Start the service.
stop — Stop the service.

Operations

void start(string name, **Ice::Communicator** communicator, **Ice::StringSeq** args)

Start the service. The given communicator is created by the IceBox::ServiceManager for use by the service. This communicator may also be used by other services, depending on the service configuration.

The IceBox::ServiceManager owns this communicator, and is responsible for destroying it.

Parameters

`name` — The service's name, as determined by the configuration.
`communicator` — A communicator for use by the service.
`args` — The service arguments that were not converted into properties.

Exceptions

IceBox::FailureException — Raised if start failed.

void stop()

Stop the service.

Copyright 2016, ZeroC, Inc.

IceBox-ServiceManager

IceBox::ServiceManager

Overview

interface ServiceManager

Administers a set of IceBox::Service instances.

See Also

- IceBox::Service

Operation Index

getSliceChecksums — Returns the checksums for the IceBox Slice definitions.
startService — Start an individual service.
stopService — Stop an individual service.
addObserver — Registers a new observer with the ServiceManager.
shutdown — Shut down all services.

Operations

["nonmutating"] **Ice::SliceChecksumDict** getSliceChecksums()

Returns the checksums for the IceBox Slice definitions.

Return Value

A dictionary mapping Slice type ids to their checksums.

void startService(string service) throws **IceBox::AlreadyStartedException**, **IceBox::NoSuchServiceException**

Start an individual service.

Parameters

`service` — The service name.

void stopService(string service) throws **IceBox::AlreadyStoppedException**, **IceBox::NoSuchServiceException**

Stop an individual service.

Parameters

`service` — The service name.

void addObserver(**IceBox::ServiceObserver*** observer)

Registers a new observer with the ServiceManager.

Parameters

`observer` — The new observer

void shutdown()

Shut down all services. This causes IceBox::Service::stop to be invoked on all configured services.

Copyright 2016, ZeroC, Inc.

IceBox-ServiceObserver

IceBox::ServiceObserver

Overview

interface ServiceObserver

An Observer interface implemented by admin clients interested in the status of services

See Also

- IceBox::ServiceManager

Operation Index

servicesStarted
servicesStopped

Operations

void servicesStarted(Ice::StringSeq services)

void servicesStopped(Ice::StringSeq services)

Copyright 2016, ZeroC, Inc.

IceGrid Slice API

IceGrid

Overview

["objc:prefix:ICEGRID"] module IceGrid

IceGrid is a server activation and deployment tool. IceGrid, simplifies the complex task of deploying applications in a heterogeneous computer network.

Class Index

AdaptiveLoadBalancingPolicy — Adaptive load balancing policy.
BoxedDistributionDescriptor — A "boxed" distribution descriptor.
BoxedString — A "boxed" string.
CommunicatorDescriptor — A communicator descriptor.
IceBoxDescriptor — An IceBox server descriptor.
LoadBalancingPolicy — A base class for load balancing policies.
OrderedLoadBalancingPolicy — Ordered load balancing policy.
RandomLoadBalancingPolicy — Random load balancing policy.
RoundRobinLoadBalancingPolicy — Round robin load balancing policy.
ServerDescriptor — An Ice server descriptor.
ServiceDescriptor — An IceBox service descriptor.

Interface Index

AdapterObserver — This interface allows applications to monitor the state of object adapters that are registered with IceGrid.
Admin — The IceGrid administrative interface.
AdminSession — Used by administrative clients to view, update, and receive observer updates from the IceGrid registry.
ApplicationObserver — The database observer interface.
FileIterator — This interface provides access to IceGrid log file contents.
FileParser — `icegridadmin` provides a FileParser object to transform XML files into ApplicationDescriptor objects.
Locator — The IceGrid locator interface provides access to the Query and Registry object of the IceGrid registry.
NodeObserver — The node observer interface.
ObjectObserver — This interface allows applications to monitor IceGrid well-known objects.
Query — The IceGrid query interface.
Registry — The IceGrid registry allows clients create sessions directly with the registry.
RegistryObserver — This interface allows applications to monitor changes the state of the registry.
RegistryPluginFacade — The RegistryPluginFacade is implemented by IceGrid and can be used by plugins and filter implementations to retrieve information from IceGrid about the well-known objects or adapters.
ReplicaGroupFilter — The ReplicaGroupFilter is used by IceGrid to filter adapters returned to the client when it resolves a filtered replica group.
Session — A session object is used by IceGrid clients to allocate and release objects.
TypeFilter — The TypeFilter is used by IceGrid to filter well-known proxies returned to the client when it searches a well-known object by type.
UserAccountMapper — A user account mapper object is used by IceGrid nodes to map session identifiers to user accounts.

Exception Index

AccessDeniedException — This exception is raised if a registry lock wasn't acquired or is already held by a session.
AdapterNotExistException — This exception is raised if an adapter does not exist.
AllocationException — This exception is raised if the allocation of an object failed.
AllocationTimeoutException — This exception is raised if the request to allocate an object times out.
ApplicationNotExistException — This exception is raised if an application does not exist.
BadSignalException — This exception is raised if an unknown signal was sent to to a server.
DeploymentException — An exception for deployment errors.
FileNotAvailableException — This exception is raised if a file is not available.
NodeNotExistException — This exception is raised if a node does not exist.
NodeUnreachableException — This exception is raised if a node could not be reached.
ObjectExistsException — This exception is raised if an object already exists.

Copyright 2016, ZeroC, Inc.

ObjectNotRegisteredException — This exception is raised if an object is not registered.
ObserverAlreadyRegisteredException — This exception is raised if an observer is already registered with the registry.
ParseException — This exception is raised if an error occurs during parsing.
PatchException — This exception is raised if a patch failed.
PermissionDeniedException — This exception is raised if a client is denied the ability to create a session with IceGrid.
RegistryNotExistException — This exception is raised if a registry does not exist.
RegistryUnreachableException — This exception is raised if a registry could not be reached.
ServerNotExistException — This exception is raised if a server does not exist.
ServerStartException — This exception is raised if a server failed to start.
ServerStopException — This exception is raised if a server failed to stop.
ServerUnreachableException — This exception is raised if a server could not be reached.
UserAccountNotFoundException — This exception is raised if a user account for a given session identifier can't be found.

Structure Index

AdapterDescriptor — An Ice object adapter descriptor.
AdapterDynamicInfo — Dynamic information about the state of an adapter.
AdapterInfo — Information about an adapter registered with the IceGrid registry.
ApplicationDescriptor — An application descriptor.
ApplicationInfo — Information about an IceGrid application.
ApplicationUpdateDescriptor — An application update descriptor to describe the updates to apply to a deployed application.
ApplicationUpdateInfo — Information about updates to an IceGrid application.
DbEnvDescriptor — A Freeze database environment descriptor.
DistributionDescriptor — A distribution descriptor defines an IcePatch2 server and the directories to retrieve from the patch server.
LoadInfo — Information about the load of a node.
NodeDescriptor — A node descriptor.
NodeDynamicInfo — Dynamic information about the state of a node.
NodeInfo — Information about an IceGrid node.
NodeUpdateDescriptor — A node update descriptor to describe the updates to apply to a node of a deployed application.
ObjectDescriptor — An Ice object descriptor.
ObjectInfo — Information about an Ice object.
PropertyDescriptor — Property descriptor.
PropertySetDescriptor — A property set descriptor.
RegistryInfo — Information about an IceGrid registry replica.
ReplicaGroupDescriptor — A replica group descriptor.
ServerDynamicInfo — Dynamic information about the state of a server.
ServerInfo — Information about a server managed by an IceGrid node.
ServerInstanceDescriptor — A server template instance descriptor.
ServiceInstanceDescriptor — A service template instance descriptor.
TemplateDescriptor — A template descriptor for server or service templates.

Sequence Index

AdapterDescriptorSeq — A sequence of adapter descriptors.
AdapterDynamicInfoSeq — A sequence of adapter dynamic information structures.
AdapterInfoSeq — A sequence of adapter information structures.
ApplicationDescriptorSeq — A sequence of application descriptors.
ApplicationInfoSeq — A sequence of ApplicationInfo structures.
DbEnvDescriptorSeq — A sequence of database environment descriptors.
NodeDynamicInfoSeq — A sequence of node dynamic information structures.
NodeUpdateDescriptorSeq — A sequence of node update descriptors.
ObjectDescriptorSeq — A sequence of object descriptors.
ObjectInfoSeq — A sequence of object information structures.
PropertyDescriptorSeq — A sequence of property descriptors.
RegistryInfoSeq — A sequence of RegistryInfo structures.
ReplicaGroupDescriptorSeq — A sequence of replica groups.
ServerDescriptorSeq — A sequence of server descriptors.
ServerDynamicInfoSeq — A sequence of server dynamic information structures.
ServerInstanceDescriptorSeq — A sequence of server instance descriptors.
ServiceDescriptorSeq — A sequence of service descriptors.
ServiceInstanceDescriptorSeq — A sequence of service instance descriptors.

Dictionary Index

NodeDescriptorDict — Mapping of node name to node descriptor.

Copyright 2016, ZeroC, Inc.

PropertySetDescriptorDict — A mapping of property set name to property set descriptor.
StringObjectProxyDict — A dictionary of proxies.
StringStringDict — A mapping of string to string.
TemplateDescriptorDict — A mapping of template identifier to template descriptor.

Enumeration Index

LoadSample — Determines which load sampling interval to use.
ServerState — An enumeration representing the state of the server.

Sequences

["java:type:java.util.LinkedList<AdapterDescriptor>"] sequence<AdapterDescriptor> AdapterDescriptorSeq

A sequence of adapter descriptors.

Used By

- CommunicatorDescriptor::adapters

["java:type:java.util.LinkedList<AdapterDynamicInfo>"] sequence<AdapterDynamicInfo> AdapterDynamicInfoSeq

A sequence of adapter dynamic information structures.

Used By

- NodeDynamicInfo::adapters

sequence<AdapterInfo> AdapterInfoSeq

A sequence of adapter information structures.

Used By

- AdapterObserver::adapterInit
- Admin::getAdapterInfo
- RegistryPluginFacade::getAdapterInfo

["java:type:java.util.LinkedList<ApplicationDescriptor>"] sequence<ApplicationDescriptor> ApplicationDescriptorSeq

A sequence of application descriptors.

["java:type:java.util.LinkedList<ApplicationInfo>"] sequence<ApplicationInfo> ApplicationInfoSeq

A sequence of ApplicationInfo structures.

Used By

- ApplicationObserver::applicationInit

["java:type:java.util.LinkedList<DbEnvDescriptor>"] sequence<DbEnvDescriptor> DbEnvDescriptorSeq

A sequence of database environment descriptors.

Used By

- CommunicatorDescriptor::dbEnvs

Copyright 2016, ZeroC, Inc.

sequence<NodeDynamicInfo> NodeDynamicInfoSeq

A sequence of node dynamic information structures.

Used By

- NodeObserver::nodeInit

["java:type:java.util.LinkedList<NodeUpdateDescriptor>"] sequence<NodeUpdateDescriptor> NodeUpdateDescriptorSeq

A sequence of node update descriptors.

Used By

- ApplicationUpdateDescriptor::nodes

["java:type:java.util.LinkedList<ObjectDescriptor>"] sequence<ObjectDescriptor> ObjectDescriptorSeq

A sequence of object descriptors.

Used By

- AdapterDescriptor::allocatables
- AdapterDescriptor::objects
- ReplicaGroupDescriptor::objects

sequence<ObjectInfo> ObjectInfoSeq

A sequence of object information structures.

Used By

- Admin::getAllObjectInfos
- Admin::getObjectInfosByType
- ObjectObserver::objectInit

["java:type:java.util.LinkedList<PropertyDescriptor>"] sequence<PropertyDescriptor> PropertyDescriptorSeq

A sequence of property descriptors.

Used By

- DbEnvDescriptor::properties
- PropertySetDescriptor::properties

sequence<RegistryInfo> RegistryInfoSeq

A sequence of RegistryInfo structures.

Used By

- RegistryObserver::registryInit

["java:type:java.util.LinkedList<ReplicaGroupDescriptor>"] sequence<ReplicaGroupDescriptor> ReplicaGroupDescriptorSeq

A sequence of replica groups.

Used By

Copyright 2016, ZeroC, Inc.

- ApplicationDescriptor::replicaGroups
- ApplicationUpdateDescriptor::replicaGroups

["java:type:java.util.LinkedList<ServerDescriptor>"] sequence<**ServerDescriptor**> ServerDescriptorSeq

A sequence of server descriptors.

Used By

- NodeDescriptor::servers
- NodeUpdateDescriptor::servers

["java:type:java.util.LinkedList<ServerDynamicInfo>"] sequence<**ServerDynamicInfo**> ServerDynamicInfoSeq

A sequence of server dynamic information structures.

Used By

- NodeDynamicInfo::servers

["java:type:java.util.LinkedList<ServerInstanceDescriptor>"] sequence<**ServerInstanceDescriptor**> ServerInstanceDescriptorSeq

A sequence of server instance descriptors.

Used By

- NodeDescriptor::serverInstances
- NodeUpdateDescriptor::serverInstances

["java:type:java.util.LinkedList<ServiceDescriptor>"] sequence<**ServiceDescriptor**> ServiceDescriptorSeq

A sequence of service descriptors.

["java:type:java.util.LinkedList<ServiceInstanceDescriptor>"] sequence<**ServiceInstanceDescriptor**> ServiceInstanceDescriptorSeq

A sequence of service instance descriptors.

Used By

- IceBoxDescriptor::services

Dictionaries

dictionary<string, **NodeDescriptor**> NodeDescriptorDict

Mapping of node name to node descriptor.

Used By

- ApplicationDescriptor::nodes

dictionary<string, **PropertySetDescriptor**> PropertySetDescriptorDict

A mapping of property set name to property set descriptor.

Used By

Copyright 2016, ZeroC, Inc.

- ApplicationDescriptor::propertySets
- ApplicationUpdateDescriptor::propertySets
- NodeDescriptor::propertySets
- NodeUpdateDescriptor::propertySets
- ServerInstanceDescriptor::servicePropertySets

dictionary<string, Object*> StringObjectProxyDict

A dictionary of proxies.

dictionary<string, string> StringStringDict

A mapping of string to string.

Used By

- ApplicationDescriptor::variables
- ApplicationUpdateDescriptor::variables
- NodeDescriptor::variables
- NodeUpdateDescriptor::variables
- ServerInstanceDescriptor::parameterValues
- ServiceInstanceDescriptor::parameterValues
- TemplateDescriptor::parameterDefaults

dictionary<string, **TemplateDescriptor**> TemplateDescriptorDict

A mapping of template identifier to template descriptor.

Used By

- ApplicationDescriptor::serverTemplates
- ApplicationDescriptor::serviceTemplates
- ApplicationUpdateDescriptor::serverTemplates
- ApplicationUpdateDescriptor::serviceTemplates

Copyright 2016, ZeroC, Inc.

IceGrid-AccessDeniedException

IceGrid::AccessDeniedException

Overview

exception AccessDeniedException

This exception is raised if a registry lock wasn't acquired or is already held by a session.

Used By

- IceGrid::Admin::addApplication
- IceGrid::Admin::instantiateServer
- IceGrid::Admin::removeApplication
- IceGrid::Admin::syncApplicationWithoutRestart
- IceGrid::Admin::syncApplication
- IceGrid::Admin::updateApplicationWithoutRestart
- IceGrid::Admin::updateApplication
- IceGrid::AdminSession::finishUpdate
- IceGrid::AdminSession::startUpdate

Data Member Index

lockUserId — The id of the user holding the lock (if any).

Data Members

string lockUserId;

The id of the user holding the lock (if any).

Copyright 2016, ZeroC, Inc.

IceGrid-AdapterDescriptor

IceGrid::AdapterDescriptor

Overview

["cpp:comparable"] struct AdapterDescriptor

An Ice object adapter descriptor.

Used By

- IceGrid::AdapterDescriptorSeq

Data Member Index

name — The object adapter name.
description — The description of this object adapter.
id — The object adapter id.
replicaGroupId — The replica id of this adapter.
priority — The adapter priority.
registerProcess — Flag to specify if the object adapter will register a process object.
serverLifetime — If true the lifetime of this object adapter is the same of the server lifetime.
objects — The well-known object descriptors associated with this object adapter.
allocatables — The allocatable object descriptors associated with this object adapter.

Data Members

string name;

The object adapter name.

string description;

The description of this object adapter.

string id;

The object adapter id.

string replicaGroupId;

The replica id of this adapter.

string priority;

The adapter priority. This is eventually used when the adapter is member of a replica group to sort the adapter endpoints by priority.

bool registerProcess;

Flag to specify if the object adapter will register a process object.

bool serverLifetime;

If true the lifetime of this object adapter is the same of the server lifetime. This information is used by the IceGrid node to figure out the server state: the server is active only if all its "server lifetime" adapters are active.

IceGrid::ObjectDescriptorSeq objects;

The well-known object descriptors associated with this object adapter.

Copyright 2016, ZeroC, Inc.

IceGrid::ObjectDescriptorSeq allocatables;

The allocatable object descriptors associated with this object adapter.

Copyright 2016, ZeroC, Inc.

IceGrid-AdapterDynamicInfo

IceGrid::AdapterDynamicInfo

Overview

struct AdapterDynamicInfo

Dynamic information about the state of an adapter.

Used By

- IceGrid::AdapterDynamicInfoSeq
- IceGrid::NodeObserver::updateAdapter

Data Member Index

id — The id of the adapter.
proxy — The direct proxy containing the adapter endpoints.

Data Members

string id;

The id of the adapter.

Object* proxy;

The direct proxy containing the adapter endpoints.

Copyright 2016, ZeroC, Inc.

IceGrid-AdapterInfo

IceGrid::AdapterInfo

Overview

["cpp:comparable"] struct AdapterInfo

Information about an adapter registered with the IceGrid registry.

Used By

- IceGrid::AdapterInfoSeq
- IceGrid::AdapterObserver::adapterAdded
- IceGrid::AdapterObserver::adapterUpdated

Data Member Index

id — The id of the adapter.
proxy — A dummy direct proxy that contains the adapter endpoints.
replicaGroupId — The replica group id of the object adapter, or empty if the adapter doesn't belong to a replica group.

Data Members

string id;

The id of the adapter.

Object* proxy;

A dummy direct proxy that contains the adapter endpoints.

string replicaGroupId;

The replica group id of the object adapter, or empty if the adapter doesn't belong to a replica group.

Copyright 2016, ZeroC, Inc.

IceGrid-AdapterNotExistException

IceGrid::AdapterNotExistException

Overview

exception AdapterNotExistException

This exception is raised if an adapter does not exist.

Used By

- IceGrid::Admin::getAdapterInfo
- IceGrid::Admin::removeAdapter
- IceGrid::RegistryPluginFacade::getAdapterApplication
- IceGrid::RegistryPluginFacade::getAdapterInfo
- IceGrid::RegistryPluginFacade::getAdapterNode
- IceGrid::RegistryPluginFacade::getAdapterServer
- IceGrid::RegistryPluginFacade::getPropertyForAdapter

Data Member Index

id — The id of the object adapter.

Data Members

string id;

The id of the object adapter.

Copyright 2016, ZeroC, Inc.

IceGrid-AdapterObserver

IceGrid::AdapterObserver

Overview

interface AdapterObserver

This interface allows applications to monitor the state of object adapters that are registered with IceGrid.

Operation Index

adapterInit — `adapterInit` is called after registration of an observer to indicate the state of the registry.
adapterAdded — The `adapterAdded` operation is called to notify an observer when a dynamically-registered adapter was added.
adapterUpdated — The adapterUpdated operation is called to notify an observer when a dynamically-registered adapter was updated.
adapterRemoved — The adapterRemoved operation is called to notify an observer when a dynamically-registered adapter was removed.

Operations

void adapterInit(**IceGrid::AdapterInfoSeq** adpts)

`adapterInit` is called after registration of an observer to indicate the state of the registry.

Parameters

`adpts` — The adapters that were dynamically registered with the registry (not through the deployment mechanism).

void adapterAdded(**IceGrid::AdapterInfo** info)

The `adapterAdded` operation is called to notify an observer when a dynamically-registered adapter was added.

Parameters

`info` — The details of the new adapter.

void adapterUpdated(**IceGrid::AdapterInfo** info)

The adapterUpdated operation is called to notify an observer when a dynamically-registered adapter was updated.

Parameters

`info` — The details of the updated adapter.

void adapterRemoved(string id)

The adapterRemoved operation is called to notify an observer when a dynamically-registered adapter was removed.

Parameters

`id` — The ID of the removed adapter.

Copyright 2016, ZeroC, Inc.

IceGrid-AdaptiveLoadBalancingPolicy

IceGrid::AdaptiveLoadBalancingPolicy

Overview

class AdaptiveLoadBalancingPolicy extends IceGrid::LoadBalancingPolicy

Adaptive load balancing policy.

Data Member Index

loadSample — The load sample to use for the load balancing.

Data Members

string loadSample;

The load sample to use for the load balancing. The allowed values for this attribute are "1", "5" and "15", representing respectively the load average over the past minute, the past 5 minutes and the past 15 minutes.

Copyright 2016, ZeroC, Inc.

IceGrid-Admin

IceGrid::Admin

Overview

interface Admin

The IceGrid administrative interface.

Allowing access to this interface is a security risk! Please see the IceGrid documentation for further information.

Operation Index

addApplication — Add an application to IceGrid.
syncApplication — Synchronize a deployed application with the given application descriptor.
updateApplication — Update a deployed application with the given update application descriptor.
syncApplicationWithoutRestart — Synchronize a deployed application with the given application descriptor.
updateApplicationWithoutRestart — Update a deployed application with the given update application descriptor only if no server restarts are necessary for the update of the application.
removeApplication — Remove an application from IceGrid.
instantiateServer — Instantiate a server template from an application on the given node.
patchApplication — Patch the given application data.
getApplicationInfo — Get an application descriptor.
getDefaultApplicationDescriptor — Get the default application descriptor.
getAllApplicationNames — Get all the IceGrid applications currently registered.
getServerInfo — Get the server information for the server with the given id.
getServerState — Get a server's state.
getServerPid — Get a server's system process id.
getServerAdminCategory — Get the category for server admin objects.
getServerAdmin — Get a proxy to the server's admin object.
enableServer — Enable or disable a server.
isServerEnabled — Check if the server is enabled or disabled.
startServer — Start a server and wait for its activation.
stopServer — Stop a server.
patchServer — Patch a server.
sendSignal — Send signal to a server.
getAllServerIds — Get all the server ids registered with IceGrid.
getAdapterInfo — Get the adapter information for the replica group or adapter with the given id.
removeAdapter — Remove the adapter with the given id.
getAllAdapterIds — Get all the adapter ids registered with IceGrid.
addObject — Add an object to the object registry.
updateObject — Update an object in the object registry.
addObjectWithType — Add an object to the object registry and explicitly specify its type.
removeObject — Remove an object from the object registry.
getObjectInfo — Get the object info for the object with the given identity.
getObjectInfosByType — Get the object info of all the registered objects with the given type.
getAllObjectInfos — Get the object info of all the registered objects whose stringified identities match the given expression.
pingNode — Ping an IceGrid node to see if it is active.
getNodeLoad — Get the load averages of the node.
getNodeInfo — Get the node information for the node with the given name.
getNodeAdmin — Get a proxy to the IceGrid node's admin object.
getNodeProcessorSocketCount — Get the number of physical processor sockets for the machine running the node with the given name.
shutdownNode — Shutdown an IceGrid node.
getNodeHostname — Get the hostname of this node.
getAllNodeNames — Get all the IceGrid nodes currently registered.
pingRegistry — Ping an IceGrid registry to see if it is active.
getRegistryInfo — Get the registry information for the registry with the given name.
getRegistryAdmin — Get a proxy to the IceGrid registry's admin object.
shutdownRegistry — Shutdown an IceGrid registry.
getAllRegistryNames — Get all the IceGrid registries currently registered.
shutdown — Shut down the IceGrid registry.
getSliceChecksums — Returns the checksums for the IceGrid Slice definitions.

Copyright 2016, ZeroC, Inc.

Operations

void addApplication(IceGrid::ApplicationDescriptor descriptor) throws IceGrid::AccessDeniedException, IceGrid::DeploymentException

Add an application to IceGrid.

Parameters

`descriptor` — The application descriptor.

Exceptions

IceGrid::AccessDeniedException — Raised if the session doesn't hold the exclusive lock or if another session is holding the lock.
IceGrid::DeploymentException — Raised if application deployment failed.

void syncApplication(IceGrid::ApplicationDescriptor descriptor) throws IceGrid::AccessDeniedException, IceGrid::DeploymentException, IceGrid::ApplicationNotExistException

Synchronize a deployed application with the given application descriptor. This operation will replace the current descriptor with this new descriptor.

Parameters

`descriptor` — The application descriptor.

Exceptions

IceGrid::AccessDeniedException — Raised if the session doesn't hold the exclusive lock or if another session is holding the lock.
IceGrid::DeploymentException — Raised if application deployment failed.
IceGrid::ApplicationNotExistException — Raised if the application doesn't exist.

void updateApplication(IceGrid::ApplicationUpdateDescriptor descriptor) throws IceGrid::AccessDeniedException, IceGrid::DeploymentException, IceGrid::ApplicationNotExistException

Update a deployed application with the given update application descriptor.

Parameters

`descriptor` — The update descriptor.

Exceptions

IceGrid::AccessDeniedException — Raised if the session doesn't hold the exclusive lock or if another session is holding the lock.
IceGrid::DeploymentException — Raised if application deployment failed.
IceGrid::ApplicationNotExistException — Raised if the application doesn't exist.

void syncApplicationWithoutRestart(IceGrid::ApplicationDescriptor descriptor) throws IceGrid::AccessDeniedException, IceGrid::DeploymentException, IceGrid::ApplicationNotExistException

Synchronize a deployed application with the given application descriptor. This operation will replace the current descriptor with this new descriptor only if no server restarts are necessary for the update of the application. If some servers need to be restarted, the synchronization is rejected with a DeploymentException.

Parameters

`descriptor` — The application descriptor.

Exceptions

IceGrid::AccessDeniedException — Raised if the session doesn't hold the exclusive lock or if another session is holding the lock.
IceGrid::DeploymentException — Raised if application deployment failed.
IceGrid::ApplicationNotExistException — Raised if the application doesn't exist.

void updateApplicationWithoutRestart(IceGrid::ApplicationUpdateDescriptor descriptor) throws IceGrid::AccessDe

Copyright 2016, ZeroC, Inc.

niedException, **IceGrid::DeploymentException**, **IceGrid::ApplicationNotExistException**

Update a deployed application with the given update application descriptor only if no server restarts are necessary for the update of the application. If some servers need to be restarted, the synchronization is rejected with a DeploymentException.

Parameters

`descriptor` — The update descriptor.

Exceptions

IceGrid::AccessDeniedException — Raised if the session doesn't hold the exclusive lock or if another session is holding the lock.
IceGrid::DeploymentException — Raised if application deployment failed.
IceGrid::ApplicationNotExistException — Raised if the application doesn't exist.

void removeApplication(string name) throws IceGrid::AccessDeniedException, IceGrid::DeploymentException, IceGrid::ApplicationNotExistException

Remove an application from IceGrid.

Parameters

`name` — The application name.

Exceptions

IceGrid::AccessDeniedException — Raised if the session doesn't hold the exclusive lock or if another session is holding the lock.
IceGrid::ApplicationNotExistException — Raised if the application doesn't exist.

void instantiateServer(string application, string node, IceGrid::ServerInstanceDescriptor desc) throws IceGrid::AccessDeniedException, IceGrid::ApplicationNotExistException, IceGrid::DeploymentException

Instantiate a server template from an application on the given node.

Parameters

`application` — The application name.
`node` — The name of the node where the server will be deployed.
`desc` — The descriptor of the server instance to deploy.

Exceptions

IceGrid::AccessDeniedException — Raised if the session doesn't hold the exclusive lock or if another session is holding the lock.
IceGrid::DeploymentException — Raised if server instantiation failed.
IceGrid::ApplicationNotExistException — Raised if the application doesn't exist.

["amd"] void patchApplication(string name, bool shutdown) throws IceGrid::ApplicationNotExistException, IceGrid::PatchException

Patch the given application data.

Parameters

`name` — The application name.
`shutdown` — If true, the servers depending on the data to patch will be shut down if necessary.

Exceptions

IceGrid::ApplicationNotExistException — Raised if the application doesn't exist.
IceGrid::PatchException — Raised if the patch failed.

["nonmutating"] IceGrid::ApplicationInfo getApplicationInfo(string name) throws IceGrid::ApplicationNotExistException

Get an application descriptor.

Copyright 2016, ZeroC, Inc.

Parameters

`name` — The application name.

Return Value

The application descriptor.

Exceptions

IceGrid::ApplicationNotExistException — Raised if the application doesn't exist.

["nonmutating"] IceGrid::ApplicationDescriptor getDefaultApplicationDescriptor() throws IceGrid::DeploymentException

Get the default application descriptor.

Exceptions

IceGrid::DeploymentException — Raised if the default application descriptor can't be accessed or is invalid.

["nonmutating"] Ice::StringSeq getAllApplicationNames()

Get all the IceGrid applications currently registered.

Return Value

The application names.

["nonmutating"] IceGrid::ServerInfo getServerInfo(string id) throws IceGrid::ServerNotExistException

Get the server information for the server with the given id.

Parameters

`id` — The server id.

Return Value

The server information.

Exceptions

IceGrid::ServerNotExistException — Raised if the server doesn't exist.

["nonmutating"] IceGrid::ServerState getServerState(string id) throws IceGrid::ServerNotExistException, IceGrid::NodeUnreachableException, IceGrid::DeploymentException

Get a server's state.

Parameters

`id` — The server id.

Return Value

The server state.

Exceptions

IceGrid::ServerNotExistException — Raised if the server doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.
IceGrid::DeploymentException — Raised if the server couldn't be deployed on the node.

["nonmutating"] int getServerPid(string id) throws IceGrid::ServerNotExistException, IceGrid::NodeUnreachableE

Copyright 2016, ZeroC, Inc.

xception, IceGrid::DeploymentException

Get a server's system process id. The process id is operating system dependent.

Parameters

`id` — The server id.

Return Value

The server's process id.

Exceptions

IceGrid::ServerNotExistException — Raised if the server doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.
IceGrid::DeploymentException — Raised if the server couldn't be deployed on the node.

string getServerAdminCategory()

Get the category for server admin objects. You can manufacture a server admin proxy from the admin proxy by changing its identity: use the server ID as name and the returned category as category.

Return Value

The category for server admin objects.

Object* getServerAdmin(string id) throws IceGrid::ServerNotExistException, IceGrid::NodeUnreachableException, IceGrid::DeploymentException

Get a proxy to the server's admin object.

Parameters

`id` — The server id.

Return Value

A proxy to the server's admin object

Exceptions

IceGrid::ServerNotExistException — Raised if the server doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.
IceGrid::DeploymentException — Raised if the server couldn't be deployed on the node.

void enableServer(string id, bool enabled) throws IceGrid::ServerNotExistException, IceGrid::NodeUnreachableException, IceGrid::DeploymentException

Enable or disable a server. A disabled server can't be started on demand or administratively. The enable state of the server is not persistent: if the node is shut down and restarted, the server will be enabled by default.

Parameters

`id` — The server id.
`enabled` — True to enable the server, false to disable it.

Exceptions

IceGrid::ServerNotExistException — Raised if the server doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.
IceGrid::DeploymentException — Raised if the server couldn't be deployed on the node.

["nonmutating"] bool isServerEnabled(string id) throws IceGrid::ServerNotExistException, IceGrid::NodeUnreachableException, IceGrid::DeploymentException

Copyright 2016, ZeroC, Inc.

Check if the server is enabled or disabled.

Parameters

`id` — The server id.

Exceptions

IceGrid::ServerNotExistException — Raised if the server doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.
IceGrid::DeploymentException — Raised if the server couldn't be deployed on the node.

["amd"] void startServer(string id) throws IceGrid::ServerNotExistException, IceGrid::ServerStartException, IceGrid::NodeUnreachableException, IceGrid::DeploymentException

Start a server and wait for its activation.

Parameters

`id` — The server id.

Exceptions

IceGrid::ServerNotExistException — Raised if the server doesn't exist.
IceGrid::ServerStartException — Raised if the server couldn't be started.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.
IceGrid::DeploymentException — Raised if the server couldn't be deployed on the node.

["amd"] void stopServer(string id) throws IceGrid::ServerNotExistException, IceGrid::ServerStopException, IceGrid::NodeUnreachableException, IceGrid::DeploymentException

Stop a server.

Parameters

`id` — The server id.

Exceptions

IceGrid::ServerNotExistException — Raised if the server doesn't exist.
IceGrid::ServerStopException — Raised if the server couldn't be stopped.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.
IceGrid::DeploymentException — Raised if the server couldn't be deployed on the node.

["amd"] void patchServer(string id, bool shutdown) throws IceGrid::ServerNotExistException, IceGrid::NodeUnreachableException, IceGrid::DeploymentException, IceGrid::PatchException

Patch a server.

Parameters

`id` — The server id.
`shutdown` — If true, servers depending on the data to patch will be shut down if necessary.

Exceptions

IceGrid::ServerNotExistException — Raised if the server doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.
IceGrid::DeploymentException — Raised if the server couldn't be deployed on the node.
IceGrid::PatchException — Raised if the patch failed.

void sendSignal(string id, string signal) throws IceGrid::ServerNotExistException, IceGrid::NodeUnreachableException, IceGrid::DeploymentException, IceGrid::BadSignalException

Send signal to a server.

Copyright 2016, ZeroC, Inc.

Parameters

`id` — The server id.
`signal` — The signal, for example SIGTERM or 15.

Exceptions

IceGrid::ServerNotExistException — Raised if the server doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.
IceGrid::DeploymentException — Raised if the server couldn't be deployed on the node.
IceGrid::BadSignalException — Raised if the signal is not recognized by the target server.

["nonmutating"] Ice::StringSeq getAllServerIds()

Get all the server ids registered with IceGrid.

Return Value

The server ids.

["nonmutating"] IceGrid::AdapterInfoSeq getAdapterInfo(string id) throws IceGrid::AdapterNotExistException

Get the adapter information for the replica group or adapter with the given id.

Parameters

`id` — The adapter id.

Return Value

A sequence of adapter information structures. If the given id refers to an adapter, this sequence will contain only one element. If the given id refers to a replica group, the sequence will contain the adapter information of each member of the replica group.

Exceptions

IceGrid::AdapterNotExistException — Raised if the adapter or replica group doesn't exist.

void removeAdapter(string id) throws IceGrid::AdapterNotExistException, IceGrid::DeploymentException

Remove the adapter with the given id.

Parameters

`id` — The adapter id.

Exceptions

IceGrid::AdapterNotExistException — Raised if the adapter doesn't exist.

["nonmutating"] Ice::StringSeq getAllAdapterIds()

Get all the adapter ids registered with IceGrid.

Return Value

The adapter ids.

void addObject(Object* obj) throws IceGrid::ObjectExistsException, IceGrid::DeploymentException

Add an object to the object registry. IceGrid will get the object type by calling `ice_id` on the given proxy. The object must be reachable.

Parameters

`obj` — The object to be added to the registry.

Copyright 2016, ZeroC, Inc.

Exceptions

IceGrid::ObjectExistsException — Raised if the object is already registered.
IceGrid::DeploymentException — Raised if the object can't be added. This might be raised if the invocation on the proxy to get the object type failed.

void updateObject(Object* obj) throws IceGrid::ObjectNotRegisteredException, IceGrid::DeploymentException

Update an object in the object registry. Only objects added with this interface can be updated with this operation. Objects added with deployment descriptors should be updated with the deployment mechanism.

Parameters

`obj` — The object to be updated to the registry.

Exceptions

IceGrid::ObjectNotRegisteredException — Raised if the object isn't registered with the registry.
IceGrid::DeploymentException — Raised if the object can't be updated. This might happen if the object was added with a deployment descriptor.

void addObjectWithType(Object* obj, string type) throws IceGrid::ObjectExistsException, IceGrid::DeploymentException

Add an object to the object registry and explicitly specify its type.

Parameters

`obj` — The object to be added to the registry.
`type` — The object type.

Exceptions

IceGrid::ObjectExistsException — Raised if the object is already registered.

void removeObject(Ice::Identity id) throws IceGrid::ObjectNotRegisteredException, IceGrid::DeploymentException

Remove an object from the object registry. Only objects added with this interface can be removed with this operation. Objects added with deployment descriptors should be removed with the deployment mechanism.

Parameters

`id` — The identity of the object to be removed from the registry.

Exceptions

IceGrid::ObjectNotRegisteredException — Raised if the object isn't registered with the registry.
IceGrid::DeploymentException — Raised if the object can't be removed. This might happen if the object was added with a deployment descriptor.

["nonmutating"] IceGrid::ObjectInfo getObjectInfo(Ice::Identity id) throws IceGrid::ObjectNotRegisteredException

Get the object info for the object with the given identity.

Parameters

`id` — The identity of the object.

Return Value

The object info.

Exceptions

IceGrid::ObjectNotRegisteredException — Raised if the object isn't registered with the registry.

Copyright 2016, ZeroC, Inc.

["nonmutating"] IceGrid::ObjectInfoSeq getObjectInfosByType(string type)

Get the object info of all the registered objects with the given type.

Parameters

`type` — The type of the object.

Return Value

The object infos.

["nonmutating"] IceGrid::ObjectInfoSeq getAllObjectInfos(string expr)

Get the object info of all the registered objects whose stringified identities match the given expression.

Parameters

`expr` — The expression to match against the stringified identities of registered objects. The expression may contain a trailing wildcard (`*`) character.

Return Value

All the object infos with a stringified identity matching the given expression.

["nonmutating"] bool pingNode(string name) throws IceGrid::NodeNotExistException

Ping an IceGrid node to see if it is active.

Parameters

`name` — The node name.

Return Value

true if the node ping succeeded, false otherwise.

Exceptions

IceGrid::NodeNotExistException — Raised if the node doesn't exist.

["nonmutating"] IceGrid::LoadInfo getNodeLoad(string name) throws IceGrid::NodeNotExistException, IceGrid::NodeUnreachableException

Get the load averages of the node.

Parameters

`name` — The node name.

Return Value

The node load information.

Exceptions

IceGrid::NodeNotExistException — Raised if the node doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.

["nonmutating"] IceGrid::NodeInfo getNodeInfo(string name) throws IceGrid::NodeNotExistException, IceGrid::NodeUnreachableException

Get the node information for the node with the given name.

Parameters

Copyright 2016, ZeroC, Inc.

`name` — The node name.

Return Value

The node information.

Exceptions

IceGrid::NodeNotExistException — Raised if the node doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.

Object* getNodeAdmin(string name) throws IceGrid::NodeNotExistException, IceGrid::NodeUnreachableException

Get a proxy to the IceGrid node's admin object.

Parameters

`name` — The IceGrid node name

Return Value

A proxy to the IceGrid node's admin object

Exceptions

IceGrid::NodeNotExistException — Raised if the node doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.

["nonmutating"] int getNodeProcessorSocketCount(string name) throws IceGrid::NodeNotExistException, IceGrid::NodeUnreachableException

Get the number of physical processor sockets for the machine running the node with the given name. Note that this method will return 1 on operating systems where this can't be automatically determined and where the IceGrid.Node.ProcessorSocketCount property for the node is not set.

Parameters

`name` — The node name.

Return Value

The number of processor sockets or 1 if the number of sockets can't determined.

Exceptions

IceGrid::NodeNotExistException — Raised if the node doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.

void shutdownNode(string name) throws IceGrid::NodeNotExistException, IceGrid::NodeUnreachableException

Shutdown an IceGrid node.

Parameters

`name` — The node name.

Exceptions

IceGrid::NodeNotExistException — Raised if the node doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.

["nonmutating"] string getNodeHostname(string name) throws IceGrid::NodeNotExistException, IceGrid::NodeUnreachableException

Get the hostname of this node.

Copyright 2016, ZeroC, Inc.

Parameters

`name` — The node name.

Return Value

The node hostname.

Exceptions

IceGrid::NodeNotExistException — Raised if the node doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.

["nonmutating"] Ice::StringSeq getAllNodeNames()

Get all the IceGrid nodes currently registered.

Return Value

The node names.

bool pingRegistry(string name) throws IceGrid::RegistryNotExistException

Ping an IceGrid registry to see if it is active.

Parameters

`name` — The registry name.

Return Value

true if the registry ping succeeded, false otherwise.

Exceptions

IceGrid::RegistryNotExistException — Raised if the registry doesn't exist.

IceGrid::RegistryInfo getRegistryInfo(string name) throws IceGrid::RegistryNotExistException, IceGrid::RegistryUnreachableException

Get the registry information for the registry with the given name.

Parameters

`name` — The registry name.

Return Value

The registry information.

Exceptions

IceGrid::RegistryNotExistException — Raised if the registry doesn't exist.
IceGrid::RegistryUnreachableException — Raised if the registry could not be reached.

Object* getRegistryAdmin(string name) throws IceGrid::RegistryNotExistException

Get a proxy to the IceGrid registry's admin object.

Parameters

`name` — The registry name

Return Value

Copyright 2016, ZeroC, Inc.

A proxy to the IceGrid registry's admin object

Exceptions

IceGrid::RegistryNotExistException — Raised if the registry doesn't exist.

void shutdownRegistry(string name) throws IceGrid::RegistryNotExistException, IceGrid::RegistryUnreachableException

Shutdown an IceGrid registry.

Parameters

`name` — The registry name.

Exceptions

IceGrid::RegistryNotExistException — Raised if the registry doesn't exist.
IceGrid::RegistryUnreachableException — Raised if the registry could not be reached.

Ice::StringSeq getAllRegistryNames()

Get all the IceGrid registries currently registered.

Return Value

The registry names.

void shutdown()

Shut down the IceGrid registry.

["nonmutating"] Ice::SliceChecksumDict getSliceChecksums()

Returns the checksums for the IceGrid Slice definitions.

Return Value

A dictionary mapping Slice type ids to their checksums.

Copyright 2016, ZeroC, Inc.

IceGrid-AdminSession

IceGrid::AdminSession

Overview

interface AdminSession extends Glacier2::Session

Used by administrative clients to view, update, and receive observer updates from the IceGrid registry. Admin sessions are created either with the IceGrid::Registry object or the registry admin Glacier2::SessionManager object.

See Also

- IceGrid::Registry
- Glacier2::SessionManager

Operation Index

keepAlive — Keep the session alive.
getAdmin — Get the admin interface.
getAdminCallbackTemplate — Get a "template" proxy for admin callback objects.
setObservers — Set the observer proxies that receive notifications when the state of the registry or nodes changes.
setObserversByIdentity — Set the observer identities that receive notifications the state of the registry or nodes changes.
startUpdate — Acquires an exclusive lock to start updating the registry applications.
finishUpdate — Finish updating the registry and release the exclusive lock.
getReplicaName — Get the name of the registry replica hosting this session.
openServerLog — Open the given server log file for reading.
openServerStdErr — Open the given server stderr file for reading.
openServerStdOut — Open the given server stdout file for reading.
openNodeStdErr — Open the given node stderr file for reading.
openNodeStdOut — Open the given node stdout file for reading.
openRegistryStdErr — Open the given registry stderr file for reading.
openRegistryStdOut — Open the given registry stdout file for reading.

Operations

void keepAlive()

Keep the session alive. Clients should call this operation regularly to prevent the server from reaping the session.

See Also

- IceGrid::Registry::getSessionTimeout

["nonmutating"] IceGrid::Admin* getAdmin()

Get the admin interface. The admin object returned by this operation can only be accessed by the session.

Return Value

The admin interface proxy.

Object* getAdminCallbackTemplate()

Get a "template" proxy for admin callback objects. An Admin client uses this proxy to set the category of its callback objects, and the published endpoints of the object adapter hosting the admin callback objects.

Return Value

A template proxy. The returned proxy is null when the Admin session was established using Glacier2.

void setObservers(IceGrid::RegistryObserver* registryObs, IceGrid::NodeObserver* nodeObs, IceGrid::Application

Copyright 2016, ZeroC, Inc.

Observer* appObs, IceGrid::AdapterObserver* adptObs, IceGrid::ObjectObserver* objObs) throws IceGrid::ObserverAlreadyRegisteredException

Set the observer proxies that receive notifications when the state of the registry or nodes changes.

Parameters

`registryObs` — The registry observer.
`nodeObs` — The node observer.
`appObs` — The application observer.
`adptObs` — The adapter observer.
`objObs` — The object observer.

Exceptions

IceGrid::ObserverAlreadyRegisteredException — Raised if an observer is already registered with this registry.

void setObserversByIdentity(Ice::Identity registryObs, Ice::Identity nodeObs, Ice::Identity appObs, Ice::Identity adptObs, Ice::Identity objObs) throws IceGrid::ObserverAlreadyRegisteredException

Set the observer identities that receive notifications the state of the registry or nodes changes. This operation should be used by clients that are using a bidirectional connection to communicate with the session.

Parameters

`registryObs` — The registry observer identity.
`nodeObs` — The node observer identity.
`appObs` — The application observer.
`adptObs` — The adapter observer.
`objObs` — The object observer.

Exceptions

IceGrid::ObserverAlreadyRegisteredException — Raised if an observer is already registered with this registry.

int startUpdate() throws IceGrid::AccessDeniedException

Acquires an exclusive lock to start updating the registry applications.

Return Value

The current serial.

Exceptions

IceGrid::AccessDeniedException — Raised if the exclusive lock can't be acquired. This might happen if the lock is currently acquired by another session.

void finishUpdate() throws IceGrid::AccessDeniedException

Finish updating the registry and release the exclusive lock.

Exceptions

IceGrid::AccessDeniedException — Raised if the session doesn't hold the exclusive lock.

string getReplicaName()

Get the name of the registry replica hosting this session.

Return Value

The replica name of the registry.

IceGrid::FileIterator* openServerLog(string id, string path, int count) throws IceGrid::FileNotAvailableException, Ice

Copyright 2016, ZeroC, Inc.

Grid::ServerNotExistException, IceGrid::NodeUnreachableException, IceGrid::DeploymentException

Open the given server log file for reading. The file can be read with the returned file iterator.

Parameters

`id` — The server id.
`path` — The path of the log file. A log file can be opened only if it's declared in the server or service deployment descriptor.
`count` — Specifies where to start reading the file. If negative, the file is read from the begining. If 0 or positive, the file is read from the last c `ount` lines.

Return Value

An iterator to read the file.

Exceptions

IceGrid::FileNotAvailableException — Raised if the file can't be read.
IceGrid::ServerNotExistException — Raised if the server doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.
IceGrid::DeploymentException — Raised if the server couldn't be deployed on the node.

IceGrid::FileIterator* openServerStdErr(string id, int count) throws IceGrid::FileNotAvailableException, IceGrid::ServerNotExistException, IceGrid::NodeUnreachableException, IceGrid::DeploymentException

Open the given server stderr file for reading. The file can be read with the returned file iterator.

Parameters

`id` — The server id.
`count` — Specifies where to start reading the file. If negative, the file is read from the begining. If 0 or positive, the file is read from the last c `ount` lines.

Return Value

An iterator to read the file.

Exceptions

IceGrid::FileNotAvailableException — Raised if the file can't be read.
IceGrid::ServerNotExistException — Raised if the server doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.
IceGrid::DeploymentException — Raised if the server couldn't be deployed on the node.

IceGrid::FileIterator* openServerStdOut(string id, int count) throws IceGrid::FileNotAvailableException, IceGrid::ServerNotExistException, IceGrid::NodeUnreachableException, IceGrid::DeploymentException

Open the given server stdout file for reading. The file can be read with the returned file iterator.

Parameters

`id` — The server id.
`count` — Specifies where to start reading the file. If negative, the file is read from the begining. If 0 or positive, the file is read from the last c `ount` lines.

Return Value

An iterator to read the file.

Exceptions

IceGrid::FileNotAvailableException — Raised if the file can't be read.
IceGrid::ServerNotExistException — Raised if the server doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.
IceGrid::DeploymentException — Raised if the server couldn't be deployed on the node.

Copyright 2016, ZeroC, Inc.

IceGrid::FileIterator* openNodeStdErr(string name, int count) throws IceGrid::FileNotAvailableException, IceGrid::NodeNotExistException, IceGrid::NodeUnreachableException

Open the given node stderr file for reading. The file can be read with the returned file iterator.

Parameters

`name` — The node name.
`count` — Specifies where to start reading the file. If negative, the file is read from the begining. If 0 or positive, the file is read from the last c`ount` lines.

Return Value

An iterator to read the file.

Exceptions

IceGrid::FileNotAvailableException — Raised if the file can't be read.
IceGrid::NodeNotExistException — Raised if the node doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.

IceGrid::FileIterator* openNodeStdOut(string name, int count) throws IceGrid::FileNotAvailableException, IceGrid::NodeNotExistException, IceGrid::NodeUnreachableException

Open the given node stdout file for reading. The file can be read with the returned file iterator.

Parameters

`name` — The node name.
`count` — Specifies where to start reading the file. If negative, the file is read from the begining. If 0 or positive, the file is read from the last c`ount` lines.

Return Value

An iterator to read the file.

Exceptions

IceGrid::FileNotAvailableException — Raised if the file can't be read.
IceGrid::NodeNotExistException — Raised if the node doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.

IceGrid::FileIterator* openRegistryStdErr(string name, int count) throws IceGrid::FileNotAvailableException, IceGrid::RegistryNotExistException, IceGrid::RegistryUnreachableException

Open the given registry stderr file for reading. The file can be read with the returned file iterator.

Parameters

`name` — The registry name.
`count` — Specifies where to start reading the file. If negative, the file is read from the begining. If 0 or positive, the file is read from the last c`ount` lines.

Return Value

An iterator to read the file.

Exceptions

IceGrid::FileNotAvailableException — Raised if the file can't be read.
IceGrid::RegistryNotExistException — Raised if the registry doesn't exist.
IceGrid::RegistryUnreachableException — Raised if the registry could not be reached.

IceGrid::FileIterator* openRegistryStdOut(string name, int count) throws IceGrid::FileNotAvailableException, IceGrid::RegistryNotExistException, IceGrid::RegistryUnreachableException

Copyright 2016, ZeroC, Inc.

Open the given registry stdout file for reading. The file can be read with the returned file iterator.

Parameters

`name` — The registry name.
`count` — Specifies where to start reading the file. If negative, the file is read from the begining. If 0 or positive, the file is read from the last `count` lines.

Return Value

An iterator to read the file.

Exceptions

IceGrid::FileNotAvailableException — Raised if the file can't be read.
IceGrid::RegistryNotExistException — Raised if the registry doesn't exist.
IceGrid::RegistryUnreachableException — Raised if the registry could not be reached.

Copyright 2016, ZeroC, Inc.

IceGrid-AllocationException

IceGrid::AllocationException

Overview

exception AllocationException

This exception is raised if the allocation of an object failed.

Derived Exceptions

- IceGrid::AllocationTimeoutException

Used By

- IceGrid::Session::allocateObjectById
- IceGrid::Session::allocateObjectByType
- IceGrid::Session::releaseObject

Data Member Index

reason — The reason why the object couldn't be allocated.

Data Members

string reason;

The reason why the object couldn't be allocated.

Copyright 2016, ZeroC, Inc.

IceGrid-AllocationTimeoutException

IceGrid::AllocationTimeoutException

Overview

exception AllocationTimeoutException extends IceGrid::AllocationException

This exception is raised if the request to allocate an object times out.

Copyright 2016, ZeroC, Inc.

IceGrid-ApplicationDescriptor

IceGrid::ApplicationDescriptor

Overview

struct ApplicationDescriptor

An application descriptor.

Used By

- IceGrid::Admin::addApplication
- IceGrid::Admin::getDefaultApplicationDescriptor
- IceGrid::Admin::syncApplicationWithoutRestart
- IceGrid::Admin::syncApplication
- IceGrid::ApplicationDescriptorSeq
- IceGrid::ApplicationInfo::descriptor
- IceGrid::FileParser::parse

Data Member Index

name — The application name.
variables — The variables defined in the application descriptor.
replicaGroups — The replica groups.
serverTemplates — The server templates.
serviceTemplates — The service templates.
nodes — The application nodes.
distrib — The application distribution.
description — The description of this application.
propertySets — Property set descriptors.

Data Members

string name;

The application name.

["java:type:java.util.TreeMap<String, String>"] IceGrid::StringStringDict variables;

The variables defined in the application descriptor.

IceGrid::ReplicaGroupDescriptorSeq replicaGroups;

The replica groups.

IceGrid::TemplateDescriptorDict serverTemplates;

The server templates.

IceGrid::TemplateDescriptorDict serviceTemplates;

The service templates.

IceGrid::NodeDescriptorDict nodes;

The application nodes.

IceGrid::DistributionDescriptor distrib;

The application distribution.

Copyright 2016, ZeroC, Inc.

string description;

The description of this application.

IceGrid::PropertySetDescriptorDict propertySets;

Property set descriptors.

Copyright 2016, ZeroC, Inc.

IceGrid-ApplicationInfo

IceGrid::ApplicationInfo

Overview

struct ApplicationInfo

Information about an IceGrid application.

Used By

- IceGrid::Admin::getApplicationInfo
- IceGrid::ApplicationInfoSeq
- IceGrid::ApplicationObserver::applicationAdded
- IceGrid::RegistryPluginFacade::getApplicationInfo

Data Member Index

uuid — Unique application identifier.
createTime — The creation time.
createUser — The user who created the application.
updateTime — The update time.
updateUser — The user who updated the application.
revision — The application revision number.
descriptor — The application descriptor

Data Members

string uuid;

Unique application identifier.

long createTime;

The creation time.

string createUser;

The user who created the application.

long updateTime;

The update time.

string updateUser;

The user who updated the application.

int revision;

The application revision number.

IceGrid::ApplicationDescriptor descriptor;

The application descriptor

Copyright 2016, ZeroC, Inc.

IceGrid-ApplicationNotExistException

IceGrid::ApplicationNotExistException

Overview

exception ApplicationNotExistException

This exception is raised if an application does not exist.

Used By

- IceGrid::Admin::getApplicationInfo
- IceGrid::Admin::instantiateServer
- IceGrid::Admin::patchApplication
- IceGrid::Admin::removeApplication
- IceGrid::Admin::syncApplicationWithoutRestart
- IceGrid::Admin::syncApplication
- IceGrid::Admin::updateApplicationWithoutRestart
- IceGrid::Admin::updateApplication
- IceGrid::RegistryPluginFacade::getApplicationInfo

Data Member Index

name — The name of the application.

Data Members

string name;

The name of the application.

Copyright 2016, ZeroC, Inc.

IceGrid-ApplicationObserver

IceGrid::ApplicationObserver

Overview

interface ApplicationObserver

The database observer interface. Observers should implement this interface to receive information about the state of the IceGrid registry database.

Operation Index

applicationInit — `applicationInit` is called after the registration of an observer to indicate the state of the registry.
applicationAdded — The `applicationAdded` operation is called to notify an observer that an application was added.
applicationRemoved — The `applicationRemoved` operation is called to notify an observer that an application was removed.
applicationUpdated — The `applicationUpdated` operation is called to notify an observer that an application was updated.

Operations

void applicationInit(int serial, IceGrid::ApplicationInfoSeq applications)

`applicationInit` is called after the registration of an observer to indicate the state of the registry.

Parameters

`serial` — The current serial number of the registry database. This serial number allows observers to make sure that their internal state is synchronized with the registry.
`applications` — The applications currently registered with the registry.

void applicationAdded(int serial, IceGrid::ApplicationInfo desc)

The `applicationAdded` operation is called to notify an observer that an application was added.

Parameters

`serial` — The new serial number of the registry database.
`desc` — The descriptor of the new application.

void applicationRemoved(int serial, string name)

The `applicationRemoved` operation is called to notify an observer that an application was removed.

Parameters

`serial` — The new serial number of the registry database.
`name` — The name of the application that was removed.

void applicationUpdated(int serial, IceGrid::ApplicationUpdateInfo desc)

The `applicationUpdated` operation is called to notify an observer that an application was updated.

Parameters

`serial` — The new serial number of the registry database.
`desc` — The descriptor of the update.

Copyright 2016, ZeroC, Inc.

IceGrid-ApplicationUpdateDescriptor

IceGrid::ApplicationUpdateDescriptor

Overview

struct ApplicationUpdateDescriptor

An application update descriptor to describe the updates to apply to a deployed application.

Used By

- IceGrid::Admin::updateApplicationWithoutRestart
- IceGrid::Admin::updateApplication
- IceGrid::ApplicationUpdateInfo::descriptor

Data Member Index

name — The name of the application to update.
description — The updated description (or null if the description wasn't updated.)
distrib — The updated distribution application descriptor.
variables — The variables to update.
removeVariables — The variables to remove.
propertySets — The property sets to update.
removePropertySets — The property sets to remove.
replicaGroups — The replica groups to update.
removeReplicaGroups — The replica groups to remove.
serverTemplates — The server templates to update.
removeServerTemplates — The ids of the server template to remove.
serviceTemplates — The service templates to update.
removeServiceTemplates — The ids of the service template to remove.
nodes — The application nodes to update.
removeNodes — The nodes to remove.

Data Members

string name;

The name of the application to update.

IceGrid::BoxedString description;

The updated description (or null if the description wasn't updated.)

IceGrid::BoxedDistributionDescriptor distrib;

The updated distribution application descriptor.

["java:type:java.util.TreeMap<String, String>"] IceGrid::StringStringDict variables;

The variables to update.

Ice::StringSeq removeVariables;

The variables to remove.

IceGrid::PropertySetDescriptorDict propertySets;

The property sets to update.

Ice::StringSeq removePropertySets;

Copyright 2016, ZeroC, Inc.

The property sets to remove.

IceGrid::ReplicaGroupDescriptorSeq replicaGroups;

The replica groups to update.

Ice::StringSeq removeReplicaGroups;

The replica groups to remove.

IceGrid::TemplateDescriptorDict serverTemplates;

The server templates to update.

Ice::StringSeq removeServerTemplates;

The ids of the server template to remove.

IceGrid::TemplateDescriptorDict serviceTemplates;

The service templates to update.

Ice::StringSeq removeServiceTemplates;

The ids of the service template to remove.

IceGrid::NodeUpdateDescriptorSeq nodes;

The application nodes to update.

Ice::StringSeq removeNodes;

The nodes to remove.

Copyright 2016, ZeroC, Inc.

IceGrid-ApplicationUpdateInfo

IceGrid::ApplicationUpdateInfo

Overview

struct ApplicationUpdateInfo

Information about updates to an IceGrid application.

Used By

- IceGrid::ApplicationObserver::applicationUpdated

Data Member Index

updateTime — The update time.
updateUser — The user who updated the application.
revision — The application revision number.
descriptor — The update descriptor.

Data Members

long updateTime;

The update time.

string updateUser;

The user who updated the application.

int revision;

The application revision number.

IceGrid::ApplicationUpdateDescriptor descriptor;

The update descriptor.

Copyright 2016, ZeroC, Inc.

IceGrid-BadSignalException

IceGrid::BadSignalException

Overview

exception BadSignalException

This exception is raised if an unknown signal was sent to to a server.

Used By

- IceGrid::Admin::sendSignal

Data Member Index

reason — The details of the unknown signal.

Data Members

string reason;

The details of the unknown signal.

Copyright 2016, ZeroC, Inc.

IceGrid-BoxedDistributionDescriptor

IceGrid::BoxedDistributionDescriptor

Overview

class BoxedDistributionDescriptor

A "boxed" distribution descriptor.

Used By

- IceGrid::ApplicationUpdateDescriptor::distrib

Data Member Index

value — The value of the boxed distribution descriptor.

Data Members

IceGrid::DistributionDescriptor value;

The value of the boxed distribution descriptor.

Copyright 2016, ZeroC, Inc.

IceGrid-BoxedString

IceGrid::BoxedString

Overview

class BoxedString

A "boxed" string.

Used By

- IceGrid::ApplicationUpdateDescriptor::description
- IceGrid::NodeUpdateDescriptor::description
- IceGrid::NodeUpdateDescriptor::loadFactor

Data Member Index

value — The value of the boxed string.

Data Members

string value;

The value of the boxed string.

Copyright 2016, ZeroC, Inc.

IceGrid-CommunicatorDescriptor

IceGrid::CommunicatorDescriptor

Overview

class CommunicatorDescriptor

A communicator descriptor.

Derived Classes and Interfaces

- IceGrid::ServerDescriptor
- IceGrid::ServiceDescriptor

Used By

- IceGrid::TemplateDescriptor::descriptor

Data Member Index

adapters — The object adapters.
propertySet — The property set.
dbEnvs — The database environments.
logs — The path of each log file.
description — A description of this descriptor.

Data Members

IceGrid::AdapterDescriptorSeq adapters;

The object adapters.

IceGrid::PropertySetDescriptor propertySet;

The property set.

IceGrid::DbEnvDescriptorSeq dbEnvs;

The database environments.

Ice::StringSeq logs;

The path of each log file.

string description;

A description of this descriptor.

Copyright 2016, ZeroC, Inc.

IceGrid-DbEnvDescriptor

IceGrid::DbEnvDescriptor

Overview

["cpp:comparable"] struct DbEnvDescriptor

A Freeze database environment descriptor.

Used By

- IceGrid::DbEnvDescriptorSeq

Data Member Index

name — The name of the database environment.
description — The description of this database environment.
dbHome — The home of the database environment (i.e., the directory where the database files will be stored).
properties — The configuration properties of the database environment.

Data Members

string name;

The name of the database environment.

string description;

The description of this database environment.

string dbHome;

The home of the database environment (i.e., the directory where the database files will be stored). If empty, the node will provide a default database directory, otherwise the directory must exist.

IceGrid::PropertyDescriptorSeq properties;

The configuration properties of the database environment.

Copyright 2016, ZeroC, Inc.

IceGrid-DeploymentException

IceGrid::DeploymentException

Overview

exception DeploymentException

An exception for deployment errors.

Used By

- IceGrid::Admin::addApplication
- IceGrid::Admin::addObjectWithType
- IceGrid::Admin::addObject
- IceGrid::Admin::enableServer
- IceGrid::Admin::getDefaultApplicationDescriptor
- IceGrid::Admin::getServerAdmin
- IceGrid::Admin::getServerPid
- IceGrid::Admin::getServerState
- IceGrid::Admin::instantiateServer
- IceGrid::Admin::isServerEnabled
- IceGrid::Admin::patchServer
- IceGrid::Admin::removeAdapter
- IceGrid::Admin::removeApplication
- IceGrid::Admin::removeObject
- IceGrid::Admin::sendSignal
- IceGrid::Admin::startServer
- IceGrid::Admin::stopServer
- IceGrid::Admin::syncApplicationWithoutRestart
- IceGrid::Admin::syncApplication
- IceGrid::Admin::updateApplicationWithoutRestart
- IceGrid::Admin::updateApplication
- IceGrid::Admin::updateObject
- IceGrid::AdminSession::openServerLog
- IceGrid::AdminSession::openServerStdErr
- IceGrid::AdminSession::openServerStdOut

Data Member Index

reason — The reason for the failure.

Data Members

string reason;

The reason for the failure.

Copyright 2016, ZeroC, Inc.

IceGrid-DistributionDescriptor

IceGrid::DistributionDescriptor

Overview

["cpp:comparable"] struct DistributionDescriptor

A distribution descriptor defines an IcePatch2 server and the directories to retrieve from the patch server.

Used By

- IceGrid::ApplicationDescriptor::distrib
- IceGrid::BoxedDistributionDescriptor::value
- IceGrid::ServerDescriptor::distrib

Data Member Index

icepatch — The proxy of the IcePatch2 server.
directories — The source directories.

Data Members

string icepatch;

The proxy of the IcePatch2 server.

["java:type:java.util.LinkedList<String>"] Ice::StringSeq directories;

The source directories.

Copyright 2016, ZeroC, Inc.

IceGrid-FileIterator

IceGrid::FileIterator

Overview

interface FileIterator

This interface provides access to IceGrid log file contents.

Operation Index

read — Read lines from the log file.
destroy — Destroy the iterator.

Operations

bool read(int size, out **Ice::StringSeq** lines) throws **IceGrid::FileNotAvailableException**

Read lines from the log file.

Parameters

`size` — Specifies the maximum number of bytes to be received. The server will ensure that the returned message doesn't exceed the given size.
`lines` — The lines read from the file. If there was nothing to read from the file since the last call to read, an empty sequence is returned. The last line of the sequence is always incomplete (and therefore no '
n' should be added when writing the last line to the to the output device).

Return Value

True if EOF is encountered.

Exceptions

IceGrid::FileNotAvailableException — Raised if there was a problem to read lines from the file.

void destroy()

Destroy the iterator.

Copyright 2016, ZeroC, Inc.

IceGrid-FileNotAvailableException

IceGrid::FileNotAvailableException

Overview

exception FileNotAvailableException

This exception is raised if a file is not available.

Used By

- IceGrid::AdminSession::openNodeStdErr
- IceGrid::AdminSession::openNodeStdOut
- IceGrid::AdminSession::openRegistryStdErr
- IceGrid::AdminSession::openRegistryStdOut
- IceGrid::AdminSession::openServerLog
- IceGrid::AdminSession::openServerStdErr
- IceGrid::AdminSession::openServerStdOut
- IceGrid::FileIterator::read

See Also

- IceGrid::AdminSession::openServerStdOut
- IceGrid::AdminSession::openServerStdErr
- IceGrid::AdminSession::openNodeStdOut
- IceGrid::AdminSession::openNodeStdErr
- IceGrid::AdminSession::openRegistryStdOut
- IceGrid::AdminSession::openRegistryStdErr

Data Member Index

reason — The reason for the failure.

Data Members

string reason;

The reason for the failure.

Copyright 2016, ZeroC, Inc.

IceGrid-FileParser

IceGrid::FileParser

Overview

interface FileParser

`icegridadmin` provides a IceGrid::FileParser object to transform XML files into IceGrid::ApplicationDescriptor objects.

Operation Index

parse — Parse a file.

Operations

IceGrid::ApplicationDescriptor **parse(string xmlFile, IceGrid::Admin* adminProxy) throws IceGrid::ParseException**

Parse a file.

Parameters

`xmlFile` — Full pathname to the file.
`adminProxy` — An Admin proxy, used only to retrieve default templates when needed. May be null.

Return Value

The application descriptor.

Exceptions

IceGrid::ParseException — Raised if an error occurred during parsing.

Copyright 2016, ZeroC, Inc.

IceGrid-IceBoxDescriptor

IceGrid::IceBoxDescriptor

Overview

class IceBoxDescriptor extends IceGrid::ServerDescriptor

An IceBox server descriptor.

Data Member Index

services — The service instances.

Data Members

IceGrid::ServiceInstanceDescriptorSeq services;

The service instances.

Copyright 2016, ZeroC, Inc.

IceGrid-LoadBalancingPolicy

IceGrid::LoadBalancingPolicy

Overview

class LoadBalancingPolicy

A base class for load balancing policies.

Derived Classes and Interfaces

- IceGrid::AdaptiveLoadBalancingPolicy
- IceGrid::OrderedLoadBalancingPolicy
- IceGrid::RandomLoadBalancingPolicy
- IceGrid::RoundRobinLoadBalancingPolicy

Used By

- IceGrid::ReplicaGroupDescriptor::loadBalancing

Data Member Index

nReplicas — The number of replicas that will be used to gather the endpoints of a replica group.

Data Members

string nReplicas;

The number of replicas that will be used to gather the endpoints of a replica group.

Copyright 2016, ZeroC, Inc.

IceGrid-LoadInfo

IceGrid::LoadInfo

Overview

struct LoadInfo

Information about the load of a node.

Used By

- IceGrid::Admin::getNodeLoad
- IceGrid::RegistryPluginFacade::getNodeLoad

Data Member Index

avg1 — The load average over the past minute.
avg5 — The load average over the past 5 minutes.
avg15 — The load average over the past 15 minutes.

Data Members

float avg1;

The load average over the past minute.

float avg5;

The load average over the past 5 minutes.

float avg15;

The load average over the past 15 minutes.

Copyright 2016, ZeroC, Inc.

IceGrid-LoadSample

IceGrid::LoadSample

Overview

enum LoadSample

Determines which load sampling interval to use.

Used By

- Query::findObjectByTypeOnLeastLoadedNode

Enumerator Index

LoadSample1 — Sample every minute.
LoadSample5 — Sample every five minutes.
LoadSample15 — Sample every fifteen minutes.

Enumerators

LoadSample1

Sample every minute.

LoadSample5

Sample every five minutes.

LoadSample15

Sample every fifteen minutes.

Copyright 2016, ZeroC, Inc.

IceGrid-Locator

IceGrid::Locator

Overview

interface Locator extends Ice::Locator

The IceGrid locator interface provides access to the IceGrid::Query and IceGrid::Registry object of the IceGrid registry.

See Also

- IceGrid::Query
- IceGrid::Registry

Operation Index

getLocalRegistry — Get the proxy of the registry object hosted by this IceGrid registry.
getLocalQuery — Get the proxy of the query object hosted by this IceGrid registry.

Operations

IceGrid::Registry* getLocalRegistry()

Get the proxy of the registry object hosted by this IceGrid registry.

Return Value

The proxy of the registry object.

IceGrid::Query* getLocalQuery()

Get the proxy of the query object hosted by this IceGrid registry.

Return Value

The proxy of the query object.

Copyright 2016, ZeroC, Inc.

IceGrid-NodeDescriptor

IceGrid::NodeDescriptor

Overview

struct NodeDescriptor

A node descriptor.

Used By

- IceGrid::NodeDescriptorDict

Data Member Index

variables — The variables defined for the node.
serverInstances — The server instances.
servers — Servers (which are not template instances).
loadFactor — Load factor of the node.
description — The description of this node.
propertySets — Property set descriptors.

Data Members

["java:type:java.util.TreeMap<String, String>"] IceGrid::StringStringDict variables;

The variables defined for the node.

IceGrid::ServerInstanceDescriptorSeq serverInstances;

The server instances.

IceGrid::ServerDescriptorSeq servers;

Servers (which are not template instances).

string loadFactor;

Load factor of the node.

string description;

The description of this node.

IceGrid::PropertySetDescriptorDict propertySets;

Property set descriptors.

Copyright 2016, ZeroC, Inc.

IceGrid-NodeDynamicInfo

IceGrid::NodeDynamicInfo

Overview

struct NodeDynamicInfo

Dynamic information about the state of a node.

Used By

- IceGrid::NodeDynamicInfoSeq
- IceGrid::NodeObserver::nodeUp

Data Member Index

info — Some static information about the node.
servers — The dynamic information of the servers deployed on this node.
adapters — The dynamic information of the adapters deployed on this node.

Data Members

IceGrid::NodeInfo info;

Some static information about the node.

IceGrid::ServerDynamicInfoSeq servers;

The dynamic information of the servers deployed on this node.

IceGrid::AdapterDynamicInfoSeq adapters;

The dynamic information of the adapters deployed on this node.

Copyright 2016, ZeroC, Inc.

IceGrid-NodeInfo

IceGrid::NodeInfo

Overview

struct NodeInfo

Information about an IceGrid node.

Used By

- IceGrid::Admin::getNodeInfo
- IceGrid::NodeDynamicInfo::info
- IceGrid::RegistryPluginFacade::getNodeInfo

Data Member Index

name — The name of the node.
os — The operating system name.
hostname — The network name of the host running this node (as defined in uname()).
release — The operation system release level (as defined in uname()).
version — The operation system version (as defined in uname()).
machine — The machine hardware type (as defined in uname()).
nProcessors — The number of processor threads on the node.
dataDir — The path to the node data directory.

Data Members

string name;

The name of the node.

string os;

The operating system name.

string hostname;

The network name of the host running this node (as defined in uname()).

string release;

The operation system release level (as defined in uname()).

string version;

The operation system version (as defined in uname()).

string machine;

The machine hardware type (as defined in uname()).

int nProcessors;

The number of processor threads on the node. For example, nProcessors is 8 on a computer with a single quad-core processor and two HT threads per core.

string dataDir;

The path to the node data directory.

Copyright 2016, ZeroC, Inc.

IceGrid-NodeNotExistException

IceGrid::NodeNotExistException

Overview

exception NodeNotExistException

This exception is raised if a node does not exist.

Used By

- IceGrid::Admin::getNodeAdmin
- IceGrid::Admin::getNodeHostname
- IceGrid::Admin::getNodeInfo
- IceGrid::Admin::getNodeLoad
- IceGrid::Admin::getNodeProcessorSocketCount
- IceGrid::Admin::pingNode
- IceGrid::Admin::shutdownNode
- IceGrid::AdminSession::openNodeStdErr
- IceGrid::AdminSession::openNodeStdOut
- IceGrid::RegistryPluginFacade::getNodeInfo
- IceGrid::RegistryPluginFacade::getNodeLoad

Data Member Index

name — The node name.

Data Members

string name;

The node name.

Copyright 2016, ZeroC, Inc.

IceGrid-NodeObserver

IceGrid::NodeObserver

Overview

interface NodeObserver

The node observer interface. Observers should implement this interface to receive information about the state of the IceGrid nodes.

Operation Index

nodeInit — The `nodeInit` operation indicates the current state of nodes.
nodeUp — The `nodeUp` operation is called to notify an observer that a node came up.
nodeDown — The `nodeDown` operation is called to notify an observer that a node went down.
updateServer — The `updateServer` operation is called to notify an observer that the state of a server changed.
updateAdapter — The `updateAdapter` operation is called to notify an observer that the state of an adapter changed.

Operations

void nodeInit(IceGrid::NodeDynamicInfoSeq nodes)

The `nodeInit` operation indicates the current state of nodes. It is called after the registration of an observer.

Parameters

`nodes` — The current state of the nodes.

void nodeUp(IceGrid::NodeDynamicInfo node)

The `nodeUp` operation is called to notify an observer that a node came up.

Parameters

`node` — The node state.

void nodeDown(string name)

The `nodeDown` operation is called to notify an observer that a node went down.

Parameters

`name` — The node name.

void updateServer(string node, IceGrid::ServerDynamicInfo updatedInfo)

The `updateServer` operation is called to notify an observer that the state of a server changed.

Parameters

`node` — The node hosting the server.
`updatedInfo` — The new server state.

void updateAdapter(string node, IceGrid::AdapterDynamicInfo updatedInfo)

The `updateAdapter` operation is called to notify an observer that the state of an adapter changed.

Parameters

`node` — The node hosting the adapter.
`updatedInfo` — The new adapter state.

Copyright 2016, ZeroC, Inc.

Copyright 2016, ZeroC, Inc.

IceGrid-NodeUnreachableException

IceGrid::NodeUnreachableException

Overview

exception NodeUnreachableException

This exception is raised if a node could not be reached.

Used By

- IceGrid::Admin::enableServer
- IceGrid::Admin::getNodeAdmin
- IceGrid::Admin::getNodeHostname
- IceGrid::Admin::getNodeInfo
- IceGrid::Admin::getNodeLoad
- IceGrid::Admin::getNodeProcessorSocketCount
- IceGrid::Admin::getServerAdmin
- IceGrid::Admin::getServerPid
- IceGrid::Admin::getServerState
- IceGrid::Admin::isServerEnabled
- IceGrid::Admin::patchServer
- IceGrid::Admin::sendSignal
- IceGrid::Admin::shutdownNode
- IceGrid::Admin::startServer
- IceGrid::Admin::stopServer
- IceGrid::AdminSession::openNodeStdErr
- IceGrid::AdminSession::openNodeStdOut
- IceGrid::AdminSession::openServerLog
- IceGrid::AdminSession::openServerStdErr
- IceGrid::AdminSession::openServerStdOut
- IceGrid::RegistryPluginFacade::getNodeInfo
- IceGrid::RegistryPluginFacade::getNodeLoad

Data Member Index

name — The name of the node that is not reachable.
reason — The reason why the node couldn't be reached.

Data Members

string name;

The name of the node that is not reachable.

string reason;

The reason why the node couldn't be reached.

Copyright 2016, ZeroC, Inc.

IceGrid-NodeUpdateDescriptor

IceGrid::NodeUpdateDescriptor

Overview

struct NodeUpdateDescriptor

A node update descriptor to describe the updates to apply to a node of a deployed application.

Used By

- IceGrid::NodeUpdateDescriptorSeq

Data Member Index

name — The name of the node to update.
description — The updated description (or null if the description wasn't updated.)
variables — The variables to update.
removeVariables — The variables to remove.
propertySets — The property sets to update.
removePropertySets — The property sets to remove.
serverInstances — The server instances to update.
servers — The servers which are not template instances to update.
removeServers — The ids of the servers to remove.
loadFactor — The updated load factor of the node (or null if the load factor was not updated).

Data Members

string name;

The name of the node to update.

IceGrid::BoxedString description;

The updated description (or null if the description wasn't updated.)

["java:type:java.util.TreeMap<String, String>"] IceGrid::StringStringDict variables;

The variables to update.

Ice::StringSeq removeVariables;

The variables to remove.

IceGrid::PropertySetDescriptorDict propertySets;

The property sets to update.

Ice::StringSeq removePropertySets;

The property sets to remove.

IceGrid::ServerInstanceDescriptorSeq serverInstances;

The server instances to update.

IceGrid::ServerDescriptorSeq servers;

The servers which are not template instances to update.

Copyright 2016, ZeroC, Inc.

Ice::StringSeq removeServers;

The ids of the servers to remove.

IceGrid::BoxedString loadFactor;

The updated load factor of the node (or null if the load factor was not updated).

Copyright 2016, ZeroC, Inc.

IceGrid-ObjectDescriptor

IceGrid::ObjectDescriptor

Overview

struct ObjectDescriptor

An Ice object descriptor.

Used By

- IceGrid::ObjectDescriptorSeq

Data Member Index

id — The identity of the object.
type — The object type.
proxyOptions — Proxy options to use with the proxy created for this Ice object.

Data Members

Ice::Identity id;

The identity of the object.

string type;

The object type.

string proxyOptions;

Proxy options to use with the proxy created for this Ice object. If empty, the proxy will be created with the proxy options specified on the object adapter or replica group.

Copyright 2016, ZeroC, Inc.

IceGrid-ObjectExistsException

IceGrid::ObjectExistsException

Overview

exception ObjectExistsException

This exception is raised if an object already exists.

Used By

- IceGrid::Admin::addObjectWithType
- IceGrid::Admin::addObject

Data Member Index

id — The identity of the object.

Data Members

Ice::Identity id;

The identity of the object.

Copyright 2016, ZeroC, Inc.

IceGrid-ObjectInfo

IceGrid::ObjectInfo

Overview

["cpp:comparable"] struct ObjectInfo

Information about an Ice object.

Used By

- IceGrid::Admin::getObjectInfo
- IceGrid::ObjectInfoSeq
- IceGrid::ObjectObserver::objectAdded
- IceGrid::ObjectObserver::objectUpdated
- IceGrid::RegistryPluginFacade::getObjectInfo

Data Member Index

proxy — The proxy of the object.
type — The type of the object.

Data Members

Object* proxy;

The proxy of the object.

string type;

The type of the object.

Copyright 2016, ZeroC, Inc.

IceGrid-ObjectNotRegisteredException

IceGrid::ObjectNotRegisteredException

Overview

exception ObjectNotRegisteredException

This exception is raised if an object is not registered.

Used By

- IceGrid::Admin::getObjectInfo
- IceGrid::Admin::removeObject
- IceGrid::Admin::updateObject
- IceGrid::RegistryPluginFacade::getObjectInfo
- IceGrid::Session::allocateObjectById
- IceGrid::Session::releaseObject

Data Member Index

id — The identity of the object.

Data Members

Ice::Identity id;

The identity of the object.

Copyright 2016, ZeroC, Inc.

IceGrid-ObjectObserver

IceGrid::ObjectObserver

Overview

interface ObjectObserver

This interface allows applications to monitor IceGrid well-known objects.

Operation Index

objectInit — `objectInit` is called after the registration of an observer to indicate the state of the registry.
objectAdded — The `objectAdded` operation is called to notify an observer when an object was added to the IceGrid::Admin interface.
objectUpdated — `objectUpdated` is called to notify an observer when an object registered with the IceGrid::Admin interface was updated.
objectRemoved — `objectRemoved` is called to notify an observer when an object registered with the IceGrid::Admin interface was removed.

Operations

void objectInit(**IceGrid::ObjectInfoSeq** objects)

`objectInit` is called after the registration of an observer to indicate the state of the registry.

Parameters

`objects` — The objects registered with the IceGrid::Admin interface (not through the deployment mechanism).

void objectAdded(**IceGrid::ObjectInfo** info)

The `objectAdded` operation is called to notify an observer when an object was added to the IceGrid::Admin interface.

Parameters

`info` — The details of the added object.

void objectUpdated(**IceGrid::ObjectInfo** info)

`objectUpdated` is called to notify an observer when an object registered with the IceGrid::Admin interface was updated.

Parameters

`info` — The details of the updated object.

void objectRemoved(**Ice::Identity** id)

`objectRemoved` is called to notify an observer when an object registered with the IceGrid::Admin interface was removed.

Parameters

`id` — The identity of the removed object.

Copyright 2016, ZeroC, Inc.

IceGrid-ObserverAlreadyRegisteredException

IceGrid::ObserverAlreadyRegisteredException

Overview

exception ObserverAlreadyRegisteredException

This exception is raised if an observer is already registered with the registry.

Used By

- IceGrid::AdminSession::setObserversByIdentity
- IceGrid::AdminSession::setObservers

See Also

- IceGrid::AdminSession::setObservers
- IceGrid::AdminSession::setObserversByIdentity

Data Member Index

id — The identity of the observer.

Data Members

Ice::Identity id;

The identity of the observer.

Copyright 2016, ZeroC, Inc.

IceGrid-OrderedLoadBalancingPolicy

IceGrid::OrderedLoadBalancingPolicy

Overview

class OrderedLoadBalancingPolicy extends IceGrid::LoadBalancingPolicy

Ordered load balancing policy.

Copyright 2016, ZeroC, Inc.

IceGrid-ParseException

IceGrid::ParseException

Overview

exception ParseException

This exception is raised if an error occurs during parsing.

Used By

- IceGrid::FileParser::parse

Data Member Index

reason — The reason for the failure.

Data Members

string reason;

The reason for the failure.

Copyright 2016, ZeroC, Inc.

IceGrid-PatchException

IceGrid::PatchException

Overview

exception PatchException

This exception is raised if a patch failed.

Used By

- IceGrid::Admin::patchApplication
- IceGrid::Admin::patchServer

Data Member Index

reasons — The reasons why the patch failed.

Data Members

Ice::StringSeq reasons;

The reasons why the patch failed.

Copyright 2016, ZeroC, Inc.

IceGrid-PermissionDeniedException

IceGrid::PermissionDeniedException

Overview

exception PermissionDeniedException

This exception is raised if a client is denied the ability to create a session with IceGrid.

Used By

- IceGrid::Registry::createAdminSessionFromSecureConnection
- IceGrid::Registry::createAdminSession
- IceGrid::Registry::createSessionFromSecureConnection
- IceGrid::Registry::createSession

Data Member Index

reason — The reason why permission was denied.

Data Members

string reason;

The reason why permission was denied.

Copyright 2016, ZeroC, Inc.

IceGrid-PropertyDescriptor

IceGrid::PropertyDescriptor

Overview

struct PropertyDescriptor

Property descriptor.

Used By

- IceGrid::PropertyDescriptorSeq

Data Member Index

name — The name of the property.
value — The value of the property.

Data Members

string name;

The name of the property.

string value;

The value of the property.

Copyright 2016, ZeroC, Inc.

IceGrid-PropertySetDescriptor

IceGrid::PropertySetDescriptor

Overview

["cpp:comparable"] struct PropertySetDescriptor

A property set descriptor.

Used By

- IceGrid::CommunicatorDescriptor::propertySet
- IceGrid::PropertySetDescriptorDict
- IceGrid::ServerInstanceDescriptor::propertySet
- IceGrid::ServiceInstanceDescriptor::propertySet

Data Member Index

references — References to named property sets.
properties — The property set properties.

Data Members

Ice::StringSeq references;

References to named property sets.

IceGrid::PropertyDescriptorSeq properties;

The property set properties.

Copyright 2016, ZeroC, Inc.

IceGrid-Query

IceGrid::Query

Overview

interface Query

The IceGrid query interface. This interface is accessible to Ice clients who wish to look up well-known objects.

Operation Index

findObjectById — Find a well-known object by identity.
findObjectByType — Find a well-known object by type.
findObjectByTypeOnLeastLoadedNode — Find a well-known object by type on the least-loaded node.
findAllObjectsByType — Find all the well-known objects with the given type.
findAllReplicas — Find all the object replicas associated with the given proxy.

Operations

["nonmutating"] Object* findObjectById(**Ice::Identity** id)

Find a well-known object by identity.

Parameters

`id` — The identity.

Return Value

The proxy or null if no such object has been found.

["nonmutating"] Object* findObjectByType(string type)

Find a well-known object by type. If there are several objects registered for the given type, the object is randomly selected.

Parameters

`type` — The object type.

Return Value

The proxy or null, if no such object has been found.

["nonmutating"] Object* findObjectByTypeOnLeastLoadedNode(string type, **IceGrid::LoadSample** sample)

Find a well-known object by type on the least-loaded node. If the registry does not know which node hosts the object (for example, because the object was registered with a direct proxy), the registry assumes the object is hosted on a node that has a load average of 1.0.

Parameters

`type` — The object type.
`sample` — The sampling interval.

Return Value

The proxy or null, if no such object has been found.

["nonmutating"] **Ice::ObjectProxySeq** findAllObjectsByType(string type)

Find all the well-known objects with the given type.

Copyright 2016, ZeroC, Inc.

Parameters

`type` — The object type.

Return Value

The proxies or an empty sequence, if no such objects have been found.

Ice::ObjectProxySeq findAllReplicas(Object* proxy)

Find all the object replicas associated with the given proxy. If the given proxy is not an indirect proxy from a replica group, an empty sequence is returned.

Parameters

`proxy` — The object proxy.

Return Value

The proxies of each object replica or an empty sequence, if the given proxy is not from a replica group.

Copyright 2016, ZeroC, Inc.

IceGrid-RandomLoadBalancingPolicy

IceGrid::RandomLoadBalancingPolicy

Overview

class RandomLoadBalancingPolicy extends IceGrid::LoadBalancingPolicy

Random load balancing policy.

Copyright 2016, ZeroC, Inc.

IceGrid-Registry

IceGrid::Registry

Overview

interface Registry

The IceGrid registry allows clients create sessions directly with the registry.

See Also

- IceGrid::Session
- IceGrid::AdminSession

Operation Index

createSession — Create a client session.
createAdminSession — Create an administrative session.
createSessionFromSecureConnection — Create a client session from a secure connection.
createAdminSessionFromSecureConnection — Create an administrative session from a secure connection.
getSessionTimeout — Get the session timeout.
getACMTimeout — Get the value of the ACM timeout.

Operations

IceGrid::Session* createSession(string userId, string password) throws IceGrid::PermissionDeniedException

Create a client session.

Parameters

`userId` — The user id.
`password` — The password for the given user id.

Return Value

A proxy for the newly created session.

Exceptions

IceGrid::PermissionDeniedException — Raised if the password for the given user id is not correct, or if the user is not allowed access.

See Also

- IceGrid::Session
- Glacier2::PermissionsVerifier

IceGrid::AdminSession* createAdminSession(string userId, string password) throws IceGrid::PermissionDeniedException

Create an administrative session.

Parameters

`userId` — The user id.
`password` — The password for the given user id.

Return Value

A proxy for the newly created session.

Copyright 2016, ZeroC, Inc.

Exceptions

IceGrid::PermissionDeniedException — Raised if the password for the given user id is not correct, or if the user is not allowed access.

See Also

- IceGrid::Session
- Glacier2::PermissionsVerifier

IceGrid::Session* createSessionFromSecureConnection() throws IceGrid::PermissionDeniedException

Create a client session from a secure connection.

Return Value

A proxy for the newly created session.

Exceptions

IceGrid::PermissionDeniedException — Raised if the password for the given user id is not correct, or if the user is not allowed access.

See Also

- IceGrid::Session
- Glacier2::SSLPermissionsVerifier

IceGrid::AdminSession* createAdminSessionFromSecureConnection() throws IceGrid::PermissionDeniedException

Create an administrative session from a secure connection.

Return Value

A proxy for the newly created session.

Exceptions

IceGrid::PermissionDeniedException — Raised if the password for the given user id is not correct, or if the user is not allowed access.

See Also

- IceGrid::Session
- Glacier2::SSLPermissionsVerifier

["nonmutating"] int getSessionTimeout()

Get the session timeout. If a client or administrative client doesn't call the session keepAlive method in the time interval defined by this timeout, IceGrid might reap the session.

Return Value

The timeout (in seconds).

See Also

- IceGrid::Session::keepAlive
- IceGrid::AdminSession::keepAlive

["nonmutating"] int getACMTimeout()

Get the value of the ACM timeout. Clients supporting ACM connection heartbeats can enable them instead of explicitly sending keep alives requests. NOTE: This method is only available since Ice 3.6.

Return Value

The timeout (in seconds).

Copyright 2016, ZeroC, Inc.

Copyright 2016, ZeroC, Inc.

IceGrid-RegistryInfo

IceGrid::RegistryInfo

Overview

struct RegistryInfo

Information about an IceGrid registry replica.

Used By

- IceGrid::Admin::getRegistryInfo
- IceGrid::RegistryInfoSeq
- IceGrid::RegistryObserver::registryUp

Data Member Index

name — The name of the registry.
hostname — The network name of the host running this registry (as defined in uname()).

Data Members

string name;

The name of the registry.

string hostname;

The network name of the host running this registry (as defined in uname()).

Copyright 2016, ZeroC, Inc.

IceGrid-RegistryNotExistException

IceGrid::RegistryNotExistException

Overview

exception RegistryNotExistException

This exception is raised if a registry does not exist.

Used By

- IceGrid::Admin::getRegistryAdmin
- IceGrid::Admin::getRegistryInfo
- IceGrid::Admin::pingRegistry
- IceGrid::Admin::shutdownRegistry
- IceGrid::AdminSession::openRegistryStdErr
- IceGrid::AdminSession::openRegistryStdOut

Data Member Index

name — The registry name.

Data Members

string name;

The registry name.

Copyright 2016, ZeroC, Inc.

IceGrid-RegistryObserver

IceGrid::RegistryObserver

Overview

interface RegistryObserver

This interface allows applications to monitor changes the state of the registry.

Operation Index

registryInit — The `registryInit` operation is called after registration of an observer to indicate the state of the registries.
registryUp — The `nodeUp` operation is called to notify an observer that a node came up.
registryDown — The `nodeDown` operation is called to notify an observer that a node went down.

Operations

void registryInit(**IceGrid::RegistryInfoSeq** registries)

The `registryInit` operation is called after registration of an observer to indicate the state of the registries.

Parameters

`registries` — The current state of the registries.

void registryUp(**IceGrid::RegistryInfo** node)

The `nodeUp` operation is called to notify an observer that a node came up.

Parameters

`node` — The node state.

void registryDown(string name)

The `nodeDown` operation is called to notify an observer that a node went down.

Parameters

`name` — The node name.

Copyright 2016, ZeroC, Inc.

IceGrid-RegistryPluginFacade

IceGrid::RegistryPluginFacade

Overview

local interface RegistryPluginFacade

The RegistryPluginFacade is implemented by IceGrid and can be used by plugins and filter implementations to retrieve information from IceGrid about the well-known objects or adapters. It's also used to register/unregister replica group and type filters.

Operation Index

getApplicationInfo — Get an application descriptor.
getServerInfo — Get the server information for the server with the given id.
getAdapterServer — Get the ID of the server to which the given adapter belongs.
getAdapterApplication — Get the name of the application to which the given adapter belongs.
getAdapterNode — Get the name of the node to which the given adapter belongs.
getAdapterInfo — Get the adapter information for the replica group or adapter with the given id.
getObjectInfo — Get the object info for the object with the given identity.
getNodeInfo — Get the node information for the node with the given name.
getNodeLoad — Get the load averages of the node.
getPropertyForAdapter — Get the property value for the given property and adapter.
addReplicaGroupFilter — Add a replica group filter.
removeReplicaGroupFilter — Remove a replica group filter.
addTypeFilter — Add a type filter.
removeTypeFilter — Remove a type filter.

Operations

["nonmutating"] IceGrid::ApplicationInfo getApplicationInfo(string name) throws IceGrid::ApplicationNotExistException

Get an application descriptor.

Parameters

`name` — The application name.

Return Value

The application descriptor.

Exceptions

IceGrid::ApplicationNotExistException — Raised if the application doesn't exist.

["nonmutating"] IceGrid::ServerInfo getServerInfo(string id) throws IceGrid::ServerNotExistException

Get the server information for the server with the given id.

Parameters

`id` — The server id.

Return Value

The server information.

Exceptions

IceGrid::ServerNotExistException — Raised if the server doesn't exist.

Copyright 2016, ZeroC, Inc.

["nonmutating"] string getAdapterServer(string adapterId) throws IceGrid::AdapterNotExistException

Get the ID of the server to which the given adapter belongs.

Parameters

`adapterId` — The adapter ID.

Return Value

The server ID or the empty string if the given identifier is not associated to an object adapter defined with an application descriptor.

Exceptions

IceGrid::AdapterNotExistException — Raised if the adapter doesn't exist.

["nonmutating"] string getAdapterApplication(string adapterId) throws IceGrid::AdapterNotExistException

Get the name of the application to which the given adapter belongs.

Parameters

`adapterId` — The adapter ID.

Return Value

The application name or the empty string if the given identifier is not associated to a replica group or object adapter defined with an application descriptor.

Exceptions

IceGrid::AdapterNotExistException — Raised if the adapter doesn't exist.

["nonmutating"] string getAdapterNode(string adapterId) throws IceGrid::AdapterNotExistException

Get the name of the node to which the given adapter belongs.

Parameters

`adapterId` — The adapter ID.

Return Value

The node name or the empty string if the given identifier is not associated to an object adapter defined with an application descriptor.

Exceptions

IceGrid::AdapterNotExistException — Raised if the adapter doesn't exist.

["nonmutating"] IceGrid::AdapterInfoSeq getAdapterInfo(string id) throws IceGrid::AdapterNotExistException

Get the adapter information for the replica group or adapter with the given id.

Parameters

`id` — The adapter id.

Return Value

A sequence of adapter information structures. If the given id refers to an adapter, this sequence will contain only one element. If the given id refers to a replica group, the sequence will contain the adapter information of each member of the replica group.

Exceptions

IceGrid::AdapterNotExistException — Raised if the adapter or replica group doesn't exist.

Copyright 2016, ZeroC, Inc.

["nonmutating"] IceGrid::ObjectInfo getObjectInfo(Ice::Identity id) throws IceGrid::ObjectNotRegisteredException

Get the object info for the object with the given identity.

Parameters

`id` — The identity of the object.

Return Value

The object info.

Exceptions

IceGrid::ObjectNotRegisteredException — Raised if the object isn't registered with the registry.

["nonmutating"] IceGrid::NodeInfo getNodeInfo(string name) throws IceGrid::NodeNotExistException, IceGrid::NodeUnreachableException

Get the node information for the node with the given name.

Parameters

`name` — The node name.

Return Value

The node information.

Exceptions

IceGrid::NodeNotExistException — Raised if the node doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.

["nonmutating"] IceGrid::LoadInfo getNodeLoad(string name) throws IceGrid::NodeNotExistException, IceGrid::NodeUnreachableException

Get the load averages of the node.

Parameters

`name` — The node name.

Return Value

The node load information.

Exceptions

IceGrid::NodeNotExistException — Raised if the node doesn't exist.
IceGrid::NodeUnreachableException — Raised if the node could not be reached.

["nonmutating"] string getPropertyForAdapter(string adapterId, string name) throws IceGrid::AdapterNotExistException

Get the property value for the given property and adapter. The property is looked up in the server or service descriptor where the adapter is defined.

Parameters

`adapterId` — The adapter ID
`name` — The name of the property.

Return Value

Copyright 2016, ZeroC, Inc.

The property value.

void addReplicaGroupFilter(string id, IceGrid::ReplicaGroupFilter filter)

Add a replica group filter.

Parameters

`id` — The identifier of the filter. This identifier must match the value of the "filter" attribute specified in the replica group descriptor. To filter dynamically registered replica groups, you should use the empty filter id.
`filter` — The filter implementation.

bool removeReplicaGroupFilter(string id, IceGrid::ReplicaGroupFilter filter)

Remove a replica group filter.

Parameters

`id` — The identifier of the filter.
`filter` — The filter implementation.

Return Value

True of the filter was removed, false otherwise.

void addTypeFilter(string type, IceGrid::TypeFilter filter)

Add a type filter.

Parameters

`type` — The type to register this filter with.
`filter` — The filter implementation.

bool removeTypeFilter(string type, IceGrid::TypeFilter filter)

Remove a type filter.

Parameters

`type` — The type to register this filter with.
`filter` — The filter implementation.

Return Value

True of the filter was removed, false otherwise.

Copyright 2016, ZeroC, Inc.

IceGrid-RegistryUnreachableException

IceGrid::RegistryUnreachableException

Overview

exception RegistryUnreachableException

This exception is raised if a registry could not be reached.

Used By

- IceGrid::Admin::getRegistryInfo
- IceGrid::Admin::shutdownRegistry
- IceGrid::AdminSession::openRegistryStdErr
- IceGrid::AdminSession::openRegistryStdOut

Data Member Index

name — The name of the registry that is not reachable.
reason — The reason why the registry couldn't be reached.

Data Members

string name;

The name of the registry that is not reachable.

string reason;

The reason why the registry couldn't be reached.

Copyright 2016, ZeroC, Inc.

IceGrid-ReplicaGroupDescriptor

IceGrid::ReplicaGroupDescriptor

Overview

struct ReplicaGroupDescriptor

A replica group descriptor.

Used By

- IceGrid::ReplicaGroupDescriptorSeq

Data Member Index

id — The id of the replica group.
loadBalancing — The load balancing policy.
proxyOptions — Default options for proxies created for the replica group.
objects — The object descriptors associated with this object adapter.
description — The description of this replica group.
filter — The filter to use for this replica group.

Data Members

string id;

The id of the replica group.

IceGrid::LoadBalancingPolicy loadBalancing;

The load balancing policy.

string proxyOptions;

Default options for proxies created for the replica group.

IceGrid::ObjectDescriptorSeq objects;

The object descriptors associated with this object adapter.

string description;

The description of this replica group.

string filter;

The filter to use for this replica group.

Copyright 2016, ZeroC, Inc.

IceGrid-ReplicaGroupFilter

IceGrid::ReplicaGroupFilter

Overview

local interface ReplicaGroupFilter

The ReplicaGroupFilter is used by IceGrid to filter adapters returned to the client when it resolves a filtered replica group. IceGrid provides the list of available adapters. The implementation of this method can use the provided context and connection to filter and return the filtered set of adapters.

Used By

- IceGrid::RegistryPluginFacade::addReplicaGroupFilter
- IceGrid::RegistryPluginFacade::removeReplicaGroupFilter

Operation Index

filter — Filter the the given set of adapters.

Operations

Ice::StringSeq filter(string replicaGroupId, Ice::StringSeq adapterIds, Ice::Connection con, Ice::Context ctx)

Filter the the given set of adapters.

Parameters

`replicaGroupId` — The replica group ID.
`adapterIds` — The adpater IDs to filter.
`con` — The connection from the Ice client which is resolving the replica group endpoints.
`ctx` — The context from the Ice client which is resolving the replica group endpoints.

Return Value

The filtered adapter IDs.

Copyright 2016, ZeroC, Inc.

IceGrid-RoundRobinLoadBalancingPolicy

IceGrid::RoundRobinLoadBalancingPolicy

Overview

class RoundRobinLoadBalancingPolicy extends IceGrid::LoadBalancingPolicy

Round robin load balancing policy.

Copyright 2016, ZeroC, Inc.

IceGrid-ServerDescriptor

IceGrid::ServerDescriptor

Overview

class ServerDescriptor extends IceGrid::CommunicatorDescriptor

An Ice server descriptor.

Derived Classes and Interfaces

- IceGrid::IceBoxDescriptor

Used By

- IceGrid::ServerDescriptorSeq
- IceGrid::ServerInfo::descriptor

Data Member Index

id — The server id.
exe — The path of the server executable.
iceVersion — The Ice version used by this server.
pwd — The path to the server working directory.
options — The command line options to pass to the server executable.
envs — The server environment variables.
activation — The server activation mode (possible values are "on-demand" or "manual").
activationTimeout — The activation timeout (an integer value representing the number of seconds to wait for activation).
deactivationTimeout — The deactivation timeout (an integer value representing the number of seconds to wait for deactivation).
applicationDistrib — Specifies if the server depends on the application distribution.
distrib — The distribution descriptor.
allocatable — Specifies if the server is allocatable.
user — The user account used to run the server.

Data Members

string id;

The server id.

string exe;

The path of the server executable.

string iceVersion;

The Ice version used by this server. This is only required if backward compatibility with servers using old Ice versions is needed (otherwise the registry will assume the server is using the same Ice version). For example "3.1.1", "3.2", "3.3.0".

string pwd;

The path to the server working directory.

["java:type:java.util.LinkedList<String>"] Ice::StringSeq options;

The command line options to pass to the server executable.

["java:type:java.util.LinkedList<String>"] Ice::StringSeq envs;

The server environment variables.

Copyright 2016, ZeroC, Inc.

string activation;

The server activation mode (possible values are "on-demand" or "manual").

string activationTimeout;

The activation timeout (an integer value representing the number of seconds to wait for activation).

string deactivationTimeout;

The deactivation timeout (an integer value representing the number of seconds to wait for deactivation).

bool applicationDistrib;

Specifies if the server depends on the application distribution.

IceGrid::DistributionDescriptor distrib;

The distribution descriptor.

bool allocatable;

Specifies if the server is allocatable.

string user;

The user account used to run the server.

Copyright 2016, ZeroC, Inc.

IceGrid-ServerDynamicInfo

IceGrid::ServerDynamicInfo

Overview

struct ServerDynamicInfo

Dynamic information about the state of a server.

Used By

- IceGrid::NodeObserver::updateServer
- IceGrid::ServerDynamicInfoSeq

Data Member Index

id — The id of the server.
state — The state of the server.
pid — The process id of the server.
enabled — Indicates whether the server is enabled.

Data Members

string id;

The id of the server.

IceGrid::ServerState state;

The state of the server.

int pid;

The process id of the server.

bool enabled;

Indicates whether the server is enabled.

Copyright 2016, ZeroC, Inc.

IceGrid-ServerInfo

IceGrid::ServerInfo

Overview

struct ServerInfo

Information about a server managed by an IceGrid node.

Used By

- IceGrid::Admin::getServerInfo
- IceGrid::RegistryPluginFacade::getServerInfo

Data Member Index

application — The server application.
uuid — The application uuid.
revision — The application revision.
node — The server node.
descriptor — The server descriptor.
sessionId — The id of the session which allocated the server.

Data Members

string application;

The server application.

string uuid;

The application uuid.

int revision;

The application revision.

string node;

The server node.

IceGrid::ServerDescriptor descriptor;

The server descriptor.

string sessionId;

The id of the session which allocated the server.

Copyright 2016, ZeroC, Inc.

IceGrid-ServerInstanceDescriptor

IceGrid::ServerInstanceDescriptor

Overview

["cpp:comparable"] struct ServerInstanceDescriptor

A server template instance descriptor.

Used By

- IceGrid::Admin::instantiateServer
- IceGrid::ServerInstanceDescriptorSeq

Data Member Index

template — The template used by this instance.
parameterValues — The template parameter values.
propertySet — The property set.
servicePropertySets — The services property sets.

Data Members

string template;

The template used by this instance.

IceGrid::StringStringDict parameterValues;

The template parameter values.

IceGrid::PropertySetDescriptor propertySet;

The property set.

IceGrid::PropertySetDescriptorDict servicePropertySets;

The services property sets. It's only valid to set these property sets if the template is an IceBox server template.

Copyright 2016, ZeroC, Inc.

IceGrid-ServerNotExistException

IceGrid::ServerNotExistException

Overview

exception ServerNotExistException

This exception is raised if a server does not exist.

Used By

- IceGrid::Admin::enableServer
- IceGrid::Admin::getServerAdmin
- IceGrid::Admin::getServerInfo
- IceGrid::Admin::getServerPid
- IceGrid::Admin::getServerState
- IceGrid::Admin::isServerEnabled
- IceGrid::Admin::patchServer
- IceGrid::Admin::sendSignal
- IceGrid::Admin::startServer
- IceGrid::Admin::stopServer
- IceGrid::AdminSession::openServerLog
- IceGrid::AdminSession::openServerStdErr
- IceGrid::AdminSession::openServerStdOut
- IceGrid::RegistryPluginFacade::getServerInfo

Data Member Index

id — The identifier of the server.

Data Members

string id;

The identifier of the server.

Copyright 2016, ZeroC, Inc.

IceGrid-ServerStartException

IceGrid::ServerStartException

Overview

exception ServerStartException

This exception is raised if a server failed to start.

Used By

- IceGrid::Admin::startServer

Data Member Index

id — The identifier of the server.
reason — The reason for the failure.

Data Members

string id;

The identifier of the server.

string reason;

The reason for the failure.

Copyright 2016, ZeroC, Inc.

IceGrid-ServerState

IceGrid::ServerState

Overview

enum ServerState

An enumeration representing the state of the server.

Used By

- Admin::getServerState
- ServerDynamicInfo::state

Enumerator Index

Inactive — The server is not running.
Activating — The server is being activated and will change to the active state when the registered server object adapters are activated or to the activation timed out state if the activation timeout expires.
ActivationTimedOut — The activation timed out state indicates that the server activation timed out.
Active — The server is running.
Deactivating — The server is being deactivated.
Destroying — The server is being destroyed.
Destroyed — The server is destroyed.

Enumerators

Inactive

The server is not running.

Activating

The server is being activated and will change to the active state when the registered server object adapters are activated or to the activation timed out state if the activation timeout expires.

ActivationTimedOut

The activation timed out state indicates that the server activation timed out.

Active

The server is running.

Deactivating

The server is being deactivated.

Destroying

The server is being destroyed.

Destroyed

The server is destroyed.

Copyright 2016, ZeroC, Inc.

IceGrid-ServerStopException

IceGrid::ServerStopException

Overview

exception ServerStopException

This exception is raised if a server failed to stop.

Used By

- IceGrid::Admin::stopServer

Data Member Index

id — The identifier of the server.
reason — The reason for the failure.

Data Members

string id;

The identifier of the server.

string reason;

The reason for the failure.

Copyright 2016, ZeroC, Inc.

IceGrid-ServerUnreachableException

IceGrid::ServerUnreachableException

Overview

exception ServerUnreachableException

This exception is raised if a server could not be reached.

Data Member Index

name — The id of the server that is not reachable.
reason — The reason why the server couldn't be reached.

Data Members

string name;

The id of the server that is not reachable.

string reason;

The reason why the server couldn't be reached.

Copyright 2016, ZeroC, Inc.

IceGrid-ServiceDescriptor

IceGrid::ServiceDescriptor

Overview

class ServiceDescriptor extends IceGrid::CommunicatorDescriptor

An IceBox service descriptor.

Used By

- IceGrid::ServiceDescriptorSeq
- IceGrid::ServiceInstanceDescriptor::descriptor

Data Member Index

name — The service name.
entry — The entry point of the IceBox service.

Data Members

string name;

The service name.

string entry;

The entry point of the IceBox service.

Copyright 2016, ZeroC, Inc.

IceGrid-ServiceInstanceDescriptor

IceGrid::ServiceInstanceDescriptor

Overview

struct ServiceInstanceDescriptor

A service template instance descriptor.

Used By

- IceGrid::ServiceInstanceDescriptorSeq

Data Member Index

template — The template used by this instance.
parameterValues — The template parameter values.
descriptor — The service definition if the instance isn't a template instance (i.e.: if the template attribute is empty).
propertySet — The property set.

Data Members

string template;

The template used by this instance.

IceGrid::StringStringDict parameterValues;

The template parameter values.

IceGrid::ServiceDescriptor descriptor;

The service definition if the instance isn't a template instance (i.e.: if the template attribute is empty).

IceGrid::PropertySetDescriptor propertySet;

The property set.

Copyright 2016, ZeroC, Inc.

IceGrid-Session

IceGrid::Session

Overview

interface Session extends Glacier2::Session

A session object is used by IceGrid clients to allocate and release objects. Client sessions are either created with the IceGrid::Registry object or the registry client Glacier2::SessionManager object.

See Also

- IceGrid::Registry
- Glacier2::SessionManager

Operation Index

keepAlive — Keep the session alive.
allocateObjectById — Allocate an object.
allocateObjectByType — Allocate an object with the given type.
releaseObject — Release an object that was allocated using `allocateObjectById` or `allocateObjectByType`.
setAllocationTimeout — Set the allocation timeout.

Operations

void keepAlive()

Keep the session alive. Clients should call this operation regularly to prevent the server from reaping the session.

See Also

- IceGrid::Registry::getSessionTimeout

["amd"] Object* allocateObjectById(Ice::Identity id) throws IceGrid::ObjectNotRegisteredException, IceGrid::AllocationException

Allocate an object. Depending on the allocation timeout, this operation might hang until the object is available or until the timeout is reached.

Parameters

`id` — The identity of the object to allocate.

Return Value

The proxy of the allocated object.

Exceptions

IceGrid::ObjectNotRegisteredException — Raised if the object with the given identity is not registered with the registry.
IceGrid::AllocationException — Raised if the object can't be allocated.

See Also

- setAllocationTimeout
- releaseObject

["amd"] Object* allocateObjectByType(string type) throws IceGrid::AllocationException

Allocate an object with the given type. Depending on the allocation timeout, this operation can block until an object becomes available or until the timeout is reached.

Copyright 2016, ZeroC, Inc.

Parameters

`type` — The type of the object.

Return Value

The proxy of the allocated object.

Exceptions

IceGrid::ObjectNotRegisteredException — Raised if no objects with the given type can be allocated.
IceGrid::AllocationException — Raised if the object could not be allocated.

See Also

- setAllocationTimeout
- releaseObject

void releaseObject(Ice::Identity id) throws IceGrid::ObjectNotRegisteredException, IceGrid::AllocationException

Release an object that was allocated using `allocateObjectById` or `allocateObjectByType`.

Parameters

`id` — The identity of the object to release.

Exceptions

IceGrid::ObjectNotRegisteredException — Raised if the object with the given identity is not registered with the registry.
IceGrid::AllocationException — Raised if the given object can't be released. This might happen if the object isn't allocatable or isn't allocated by the session.

void setAllocationTimeout(int timeout)

Set the allocation timeout. If no objects are available for an allocation request, a call to `allocateObjectById` or `allocateObjectByTyp e` will block for the duration of this timeout.

Parameters

`timeout` — The timeout in milliseconds.

Copyright 2016, ZeroC, Inc.

IceGrid-TemplateDescriptor

IceGrid::TemplateDescriptor

Overview

struct TemplateDescriptor

A template descriptor for server or service templates.

Used By

- IceGrid::TemplateDescriptorDict

Data Member Index

descriptor — The template.
parameters — The parameters required to instantiate the template.
parameterDefaults — The parameters default values.

Data Members

IceGrid::CommunicatorDescriptor descriptor;

The template.

["java:type:java.util.LinkedList<String>"] Ice::StringSeq parameters;

The parameters required to instantiate the template.

IceGrid::StringStringDict parameterDefaults;

The parameters default values.

Copyright 2016, ZeroC, Inc.

IceGrid-TypeFilter

IceGrid::TypeFilter

Overview

local interface TypeFilter

The TypeFilter is used by IceGrid to filter well-known proxies returned to the client when it searches a well-known object by type. IceGrid provides the list of available proxies. The implementation of this method can use the provided context and connection to filter and return the filtered set of proxies.

Used By

- IceGrid::RegistryPluginFacade::addTypeFilter
- IceGrid::RegistryPluginFacade::removeTypeFilter

Operation Index

filter — Filter the the given set of proxies.

Operations

Ice::ObjectProxySeq filter(string type, Ice::ObjectProxySeq proxies, Ice::Connection con, Ice::Context ctx)

Filter the the given set of proxies.

Parameters

`type` — The type.
`proxies` — The proxies to filter.
`con` — The connection from the Ice client which is looking up well-known objects by type.
`ctx` — The context from the Ice client which is looking up well-known objects by type.

Return Value

The filtered proxies.

Copyright 2016, ZeroC, Inc.

IceGrid-UserAccountMapper

IceGrid::UserAccountMapper

Overview

interface UserAccountMapper

A user account mapper object is used by IceGrid nodes to map session identifiers to user accounts.

Operation Index

getUserAccount — Get the name of the user account for the given user.

Operations

string getUserAccount(string user) throws IceGrid::UserAccountNotFoundException

Get the name of the user account for the given user. This is used by IceGrid nodes to figure out the user account to use to run servers.

Parameters

`user` — The value of the server descriptor's `user` attribute. If this attribute is not defined, and the server's activation mode is `session`, the default value of `user` is the session identifier.

Return Value

The user account name.

Exceptions

IceGrid::UserAccountNotFoundException — Raised if no user account is found for the given user.

Copyright 2016, ZeroC, Inc.

IceGrid-UserAccountNotFoundException

IceGrid::UserAccountNotFoundException

Overview

exception UserAccountNotFoundException

This exception is raised if a user account for a given session identifier can't be found.

Used By

- IceGrid::UserAccountMapper::getUserAccount

Copyright 2016, ZeroC, Inc.

IceLocatorDiscovery Slice API

IceLocatorDiscovery

Overview

module IceLocatorDiscovery

IceLocatorDiscovery is an Ice plug-in that enables the discovery of IceGrid and custom locators via UDP multicast.

Interface Index

Lookup — The Ice lookup interface is implemented by Ice locator implementations and can be used by clients to find available Ice locators on the network.
LookupReply — The Ice lookup reply interface must be implemented by clients which are searching for Ice locators.

Copyright 2016, ZeroC, Inc.

IceLocatorDiscovery-Lookup

IceLocatorDiscovery::Lookup

Overview

interface Lookup

The Ice lookup interface is implemented by Ice locator implementations and can be used by clients to find available Ice locators on the network. Ice locator implementations provide a well-known `Ice/LocatorLookup' object accessible through UDP multicast. Clients typically make a multicast findLocator request to find the locator proxy.

See Also

- Ice::Locator
- IceLocatorDiscovery::LookupReply

Operation Index

findLocator — Find a locator proxy with the given instance name.

Operations

void findLocator(string instanceName, **IceLocatorDiscovery::LookupReply*** reply)

Find a locator proxy with the given instance name.

Parameters

`instanceName` — Restrict the search to Ice registries configured with the given instance name. If empty, all the available registries will reply.
`reply` — The reply object to use to send the reply.

Copyright 2016, ZeroC, Inc.

IceLocatorDiscovery-LookupReply

IceLocatorDiscovery::LookupReply

Overview

interface LookupReply

The Ice lookup reply interface must be implemented by clients which are searching for Ice locators. Ice locator implementations invoke on this interface to provide their locator proxy.

See Also

- Ice::Locator
- IceLocatorDiscovery::Lookup

Operation Index

foundLocator — This method is called by the implementation of the Lookup interface to reply to a findLocator request.

Operations

void foundLocator(**Ice::Locator*** prx)

This method is called by the implementation of the Lookup interface to reply to a findLocator request.

Parameters

`prx` — The proxy of the locator.

Copyright 2016, ZeroC, Inc.

IceMX Slice API

IceMX

Overview

["objc:prefix:ICEMX"] module IceMX

The Ice Management eXtension facility. It provides the MetricsAdmin interface for management clients to retrieve metrics from Ice applications.

Class Index

ChildInvocationMetrics — Provides information on child invocations.
CollocatedMetrics — Provides information on invocations that are collocated.
ConnectionMetrics — Provides information on the data sent and received over Ice connections.
DispatchMetrics — Provides information on servant dispatch.
InvocationMetrics — Provide measurements for proxy invocations.
Metrics — The base class for metrics.
RemoteMetrics — Provides information on invocations that are specifically sent over Ice connections.
SessionMetrics — Provides information on Glacier2 sessions.
SubscriberMetrics — Provides information on IceStorm subscribers.
ThreadMetrics — Provides information on the number of threads currently in use and their activity.
TopicMetrics — Provides information on IceStorm topics.

Interface Index

MetricsAdmin — The metrics administrative facet interface.

Exception Index

UnknownMetricsView — Raised if a metrics view cannot be found.

Structure Index

MetricsFailures — A structure to keep track of failures associated with a given metrics.

Sequence Index

MetricsFailuresSeq — A sequence of MetricsFailures.
MetricsMap — A metrics map is a sequence of metrics.

Dictionary Index

MetricsView — A metrics view is a dictionary of metrics map.
StringIntDict — A dictionnary of strings to integers.

Sequences

sequence<MetricsFailures> MetricsFailuresSeq

A sequence of MetricsFailures.

Used By

- MetricsAdmin::getMapMetricsFailures

Copyright 2016, ZeroC, Inc.

sequence<Metrics> MetricsMap

A metrics map is a sequence of metrics. We use a sequence here instead of a map because the ID of the metrics is already included in the Metrics class and using sequences of metrics objects is more efficient than using dictionaries since lookup is not necessary.

Used By

- InvocationMetrics::collocated
- InvocationMetrics::remotes
- MetricsView

Dictionaries

dictionary<string, MetricsMap> MetricsView

A metrics view is a dictionary of metrics map. The key of the dictionary is the name of the metrics map.

Used By

- MetricsAdmin::getMetricsView

dictionary<string, int> StringIntDict

A dictionnary of strings to integers.

Used By

- MetricsFailures::failures

Copyright 2016, ZeroC, Inc.

IceMX-ChildInvocationMetrics

IceMX::ChildInvocationMetrics

Overview

class ChildInvocationMetrics extends IceMX::Metrics

Provides information on child invocations. A child invocation is either remote (sent over an Ice connection) or collocated. An invocation can have multiple child invocation if it is retried. Child invocation metrics are embedded within IceMX::InvocationMetrics.

Derived Classes and Interfaces

- IceMX::CollocatedMetrics
- IceMX::RemoteMetrics

Data Member Index

size — The size of the invocation.
replySize — The size of the invocation reply.

Data Members

long size;

The size of the invocation. This corresponds to the size of the marshalled input parameters.

long replySize;

The size of the invocation reply. This corresponds to the size of the marshalled output and return parameters.

Copyright 2016, ZeroC, Inc.

IceMX-CollocatedMetrics

IceMX::CollocatedMetrics

Overview

class CollocatedMetrics extends IceMX::ChildInvocationMetrics

Provides information on invocations that are collocated. Collocated metrics are embedded within IceMX::InvocationMetrics.

Copyright 2016, ZeroC, Inc.

IceMX-ConnectionMetrics

IceMX::ConnectionMetrics

Overview

class ConnectionMetrics extends IceMX::Metrics

Provides information on the data sent and received over Ice connections.

Data Member Index

receivedBytes — The number of bytes received by the connection.
sentBytes — The number of bytes sent by the connection.

Data Members

long receivedBytes;

The number of bytes received by the connection.

long sentBytes;

The number of bytes sent by the connection.

Copyright 2016, ZeroC, Inc.

IceMX-DispatchMetrics

IceMX::DispatchMetrics

Overview

class DispatchMetrics extends IceMX::Metrics

Provides information on servant dispatch.

Data Member Index

userException — The number of dispatch that failed with a user exception.
size — The size of the dispatch.
replySize — The size of the dispatch reply.

Data Members

int userException;

The number of dispatch that failed with a user exception.

long size;

The size of the dispatch. This corresponds to the size of the marshalled input parameters.

long replySize;

The size of the dispatch reply. This corresponds to the size of the marshalled output and return parameters.

Copyright 2016, ZeroC, Inc.

IceMX-InvocationMetrics

IceMX::InvocationMetrics

Overview

class InvocationMetrics extends IceMX::Metrics

Provide measurements for proxy invocations. Proxy invocations can either be sent over the wire or be collocated.

Data Member Index

retry — The number of retries for the invocation(s).
userException — The number of invocations that failed with a user exception.
remotes — The remote invocation metrics map.
collocated — The collocated invocation metrics map.

Data Members

int retry;

The number of retries for the invocation(s).

int userException;

The number of invocations that failed with a user exception.

IceMX::MetricsMap remotes;

The remote invocation metrics map.

See Also

- IceMX::RemoteMetrics

IceMX::MetricsMap collocated;

The collocated invocation metrics map.

See Also

- IceMX::CollocatedMetrics

Copyright 2016, ZeroC, Inc.

IceMX-Metrics

IceMX::Metrics

Overview

class Metrics

The base class for metrics. A metrics object represents a collection of measurements associated to a given a system.

Derived Classes and Interfaces

- IceMX::ChildInvocationMetrics
- IceMX::ConnectionMetrics
- IceMX::DispatchMetrics
- IceMX::InvocationMetrics
- IceMX::SessionMetrics
- IceMX::SubscriberMetrics
- IceMX::ThreadMetrics
- IceMX::TopicMetrics

Used By

- IceMX::MetricsMap

Data Member Index

id — The metrics identifier.
total — The total number of objects that were observed by this metrics.
current — The current number of objects observed by this metrics.
totalLifetime — The sum of the lifetime of each observed objects.
failures — The number of failures observed.

Data Members

string id;

The metrics identifier.

long total;

The total number of objects that were observed by this metrics.

int current;

The current number of objects observed by this metrics.

long totalLifetime;

The sum of the lifetime of each observed objects. This does not include the lifetime of objects which are currently observed.

int failures;

The number of failures observed.

Copyright 2016, ZeroC, Inc.

IceMX-MetricsAdmin

IceMX::MetricsAdmin

Overview

["format:sliced"] interface MetricsAdmin

The metrics administrative facet interface. This interface allows remote administrative clients to access metrics of an application that enabled the Ice administrative facility and configured some metrics views.

Operation Index

getMetricsViewNames — Get the names of enabled and disabled metrics.
enableMetricsView — Enables a metrics view.
disableMetricsView — Disable a metrics view.
getMetricsView — Get the metrics objects for the given metrics view.
getMapMetricsFailures — Get the metrics failures associated with the given view and map.
getMetricsFailures — Get the metrics failure associated for the given metrics.

Operations

Ice::StringSeq getMetricsViewNames(out Ice::StringSeq disabledViews)

Get the names of enabled and disabled metrics.

Parameters

`disabledViews` — The names of the disabled views.

Return Value

The name of the enabled views.

void enableMetricsView(string name) throws IceMX::UnknownMetricsView

Enables a metrics view.

Parameters

`name` — The metrics view name.

Exceptions

IceMX::UnknownMetricsView — Raised if the metrics view cannot be found.

void disableMetricsView(string name) throws IceMX::UnknownMetricsView

Disable a metrics view.

Parameters

`name` — The metrics view name.

Exceptions

IceMX::UnknownMetricsView — Raised if the metrics view cannot be found.

IceMX::MetricsView getMetricsView(string view, out long timestamp) throws IceMX::UnknownMetricsView

Get the metrics objects for the given metrics view. This returns a dictionnary of metric maps for each metrics class configured with the view. The timestamp allows the client to compute averages which are not dependent of the invocation latency for this operation.

Copyright 2016, ZeroC, Inc.

Parameters

`view` — The name of the metrics view.
`timestamp` — The local time of the process when the metrics object were retrieved.

Return Value

The metrics view data.

Exceptions

IceMX::UnknownMetricsView — Raised if the metrics view cannot be found.

IceMX::MetricsFailuresSeq getMapMetricsFailures(string view, string map) throws IceMX::UnknownMetricsView

Get the metrics failures associated with the given view and map.

Parameters

`view` — The name of the metrics view.
`map` — The name of the metrics map.

Return Value

The metrics failures associated with the map.

Exceptions

IceMX::UnknownMetricsView — Raised if the metrics view cannot be found.

IceMX::MetricsFailures getMetricsFailures(string view, string map, string id) throws IceMX::UnknownMetricsView

Get the metrics failure associated for the given metrics.

Parameters

`view` — The name of the metrics view.
`map` — The name of the metrics map.
`id` — The ID of the metrics.

Return Value

The metrics failures associated with the metrics.

Exceptions

IceMX::UnknownMetricsView — Raised if the metrics view cannot be found.

Copyright 2016, ZeroC, Inc.

IceMX-MetricsFailures

IceMX::MetricsFailures

Overview

struct MetricsFailures

A structure to keep track of failures associated with a given metrics.

Used By

- IceMX::MetricsAdmin::getMetricsFailures
- IceMX::MetricsFailuresSeq

Data Member Index

id — The identifier of the metrics object associated to the failures.
failures — The failures observed for this metrics.

Data Members

string id;

The identifier of the metrics object associated to the failures.

IceMX::StringIntDict failures;

The failures observed for this metrics.

Copyright 2016, ZeroC, Inc.

IceMX-RemoteMetrics

IceMX::RemoteMetrics

Overview

class RemoteMetrics extends IceMX::ChildInvocationMetrics

Provides information on invocations that are specifically sent over Ice connections. Remote metrics are embedded within IceMX::InvocationMetrics.

Copyright 2016, ZeroC, Inc.

IceMX-SessionMetrics

IceMX::SessionMetrics

Overview

class SessionMetrics extends IceMX::Metrics

Provides information on Glacier2 sessions.

Data Member Index

forwardedClient — Number of client requests forwared.
forwardedServer — Number of server requests forwared.
routingTableSize — The size of the routing table.
queuedClient — Number of client requests queued.
queuedServer — Number of server requests queued.
overriddenClient — Number of client requests overridden.
overriddenServer — Number of server requests overridden.

Data Members

int forwardedClient;

Number of client requests forwared.

int forwardedServer;

Number of server requests forwared.

int routingTableSize;

The size of the routing table.

int queuedClient;

Number of client requests queued.

int queuedServer;

Number of server requests queued.

int overriddenClient;

Number of client requests overridden.

int overriddenServer;

Number of server requests overridden.

Copyright 2016, ZeroC, Inc.

IceMX-SubscriberMetrics

IceMX::SubscriberMetrics

Overview

class SubscriberMetrics extends IceMX::Metrics

Provides information on IceStorm subscribers.

Data Member Index

queued — Number of queued events.
outstanding — Number of outstanding events.
delivered — Number of forwarded events.

Data Members

int queued;

Number of queued events.

int outstanding;

Number of outstanding events.

long delivered;

Number of forwarded events.

Copyright 2016, ZeroC, Inc.

IceMX-ThreadMetrics

IceMX::ThreadMetrics

Overview

class ThreadMetrics extends IceMX::Metrics

Provides information on the number of threads currently in use and their activity.

Data Member Index

inUseForIO — The number of threads which are currently performing socket read or writes.
inUseForUser — The number of threads which are currently calling user code (servant dispatch, AMI callbacks, etc).
inUseForOther — The number of threads which are currently performing other activities.

Data Members

int inUseForIO;

The number of threads which are currently performing socket read or writes.

int inUseForUser;

The number of threads which are currently calling user code (servant dispatch, AMI callbacks, etc).

int inUseForOther;

The number of threads which are currently performing other activities. These are all other that are not counted with inUseForUser or inUseForIO, such as DNS lookups, garbage collection).

Copyright 2016, ZeroC, Inc.

IceMX-TopicMetrics

IceMX::TopicMetrics

Overview

class TopicMetrics extends IceMX::Metrics

Provides information on IceStorm topics.

Data Member Index

published — Number of events published on the topic by publishers.
forwarded — Number of events forwarded on the topic by IceStorm topic links.

Data Members

long published;

Number of events published on the topic by publishers.

long forwarded;

Number of events forwarded on the topic by IceStorm topic links.

Copyright 2016, ZeroC, Inc.

IceMX-UnknownMetricsView

IceMX::UnknownMetricsView

Overview

exception UnknownMetricsView

Raised if a metrics view cannot be found.

Used By

- IceMX::MetricsAdmin::disableMetricsView
- IceMX::MetricsAdmin::enableMetricsView
- IceMX::MetricsAdmin::getMapMetricsFailures
- IceMX::MetricsAdmin::getMetricsFailures
- IceMX::MetricsAdmin::getMetricsView

Copyright 2016, ZeroC, Inc.

IcePatch2 Slice API

IcePatch2

Overview

["objc:prefix:ICEPATCH2"] module IcePatch2

IcePatch can be used to update file hierarchies in a simple and efficient manner. Checksums ensure file integrity, and data is compressed before downloading.

Interface Index

FileServer — The interface that provides access to files.

Exception Index

FileAccessException — This exception is raised if a file's contents cannot be read.
FileSizeRangeException — This exception is raised if an operation tries to use a file whose size is larger than 2.1 GB.
PartitionOutOfRangeException — A `partition` argument for was not in the range 0-255.

Structure Index

FileInfo — Basic information about a single file.
LargeFileInfo — Basic information about a single file.

Sequence Index

ByteSeqSeq — A sequence of byte sequences.
FileInfoSeq — A sequence with information about many files.
LargeFileInfoSeq — A sequence with information about many files.

Sequences

sequence<Ice::ByteSeq> ByteSeqSeq

A sequence of byte sequences. Each element is the checksum for a partition.

Used By

- FileServer::getChecksumSeq

sequence<FileInfo> FileInfoSeq

A sequence with information about many files.

Used By

- FileServer::getFileInfoSeq

sequence<LargeFileInfo> LargeFileInfoSeq

A sequence with information about many files.

Used By

- FileServer::getLargeFileInfoSeq

Copyright 2016, ZeroC, Inc.

Copyright 2016, ZeroC, Inc.

IcePatch2-FileAccessException

IcePatch2::FileAccessException

Overview

exception FileAccessException

This exception is raised if a file's contents cannot be read.

Derived Exceptions

- IcePatch2::FileSizeRangeException

Used By

- IcePatch2::FileServer::getFileCompressed
- IcePatch2::FileServer::getLargeFileCompressed

Data Member Index

reason — An explanation of the reason for the failure.

Data Members

string reason;

An explanation of the reason for the failure.

Copyright 2016, ZeroC, Inc.

IcePatch2-FileInfo

IcePatch2::FileInfo

Overview

struct FileInfo

Basic information about a single file.

Used By

- IcePatch2::FileInfoSeq

Data Member Index

path — The pathname.
checksum — The SHA-1 checksum of the file.
size — The size of the compressed file in number of bytes.
executable — The executable flag.

Data Members

string path;

The pathname.

Ice::ByteSeq checksum;

The SHA-1 checksum of the file.

int size;

The size of the compressed file in number of bytes.

bool executable;

The executable flag.

Copyright 2016, ZeroC, Inc.

IcePatch2-FileServer

IcePatch2::FileServer

Overview

interface FileServer

The interface that provides access to files.

Operation Index

getFileInfoSeq — Return file information for the specified partition. *(Deprecated)*
getLargeFileInfoSeq — Returns file information for the specified partition.
getChecksumSeq — Return the checksums for all partitions.
getChecksum — Return the master checksum for all partitions.
getFileCompressed — Read the specified file. *(Deprecated)*
getLargeFileCompressed — Read the specified file.

Operations

["nonmutating"] IcePatch2::FileInfoSeq getFileInfoSeq(int partition) throws IcePatch2::PartitionOutOfRangeException, IcePatch2::FileSizeRangeException

Return file information for the specified partition.

This operation is deprecated and only present for compatibility with old Ice clients (older than version 3.6).

getFileInfoSeq() is deprecated, use getLargeFileInfoSeq() instead.

Parameters

`partition` — The partition number in the range 0-255.

Return Value

A sequence containing information about the files in the specified partition.

Exceptions

IcePatch2::PartitionOutOfRangeException — If the partition number is out of range.
IcePatch2::FileSizeRangeException — If a file is larger than 2.1GB.

["nonmutating"] IcePatch2::LargeFileInfoSeq getLargeFileInfoSeq(int partition) throws IcePatch2::PartitionOutOfRangeException

Returns file information for the specified partition.

Parameters

`partition` — The partition number in the range 0-255.

Return Value

A sequence containing information about the files in the specified partition.

Exceptions

IcePatch2::PartitionOutOfRangeException — If the partition number is out of range.

["nonmutating"] IcePatch2::ByteSeqSeq getChecksumSeq()

Return the checksums for all partitions.

Copyright 2016, ZeroC, Inc.

Return Value

A sequence containing 256 checksums. Partitions with a checksum that differs from the previous checksum for the same partition contain updated files. Partitions with a checksum that is identical to the previous checksum do not contain updated files.

["nonmutating"] Ice::ByteSeq getChecksum()

Return the master checksum for all partitions. If this checksum is the same as for a previous run, the entire file set is up-to-date.

Return Value

The master checksum for the file set.

["amd", "nonmutating", "cpp:array"] Ice::ByteSeq getFileCompressed(string path, int pos, int num) throws IcePatch2::FileAccessException, IcePatch2::FileSizeRangeException

Read the specified file. This operation may only return fewer bytes than requested in case there was an end-of-file condition.

This operation is deprecated and only present for compatibility with old Ice clients (older than version 3.6).

getFileCompressed() is deprecated, use getLargeFileCompressed() instead.

Parameters

`path` — The pathname (relative to the data directory) for the file to be read.
`pos` — The file offset at which to begin reading.
`num` — The number of bytes to be read.

Return Value

A sequence containing the compressed file contents.

Exceptions

IcePatch2::FileAccessException — If an error occurred while trying to read the file.
IcePatch2::FileSizeRangeException — If a file is larger than 2.1GB.

["amd", "nonmutating", "cpp:array"] Ice::ByteSeq getLargeFileCompressed(string path, long pos, int num) throws IcePatch2::FileAccessException

Read the specified file. This operation may only return fewer bytes than requested in case there was an end-of-file condition.

Parameters

`path` — The pathname (relative to the data directory) for the file to be read.
`pos` — The file offset at which to begin reading.
`num` — The number of bytes to be read.

Return Value

A sequence containing the compressed file contents.

Exceptions

IcePatch2::FileAccessException — If an error occurred while trying to read the file.

Copyright 2016, ZeroC, Inc.

IcePatch2-FileSizeRangeException

IcePatch2::FileSizeRangeException

Overview

exception FileSizeRangeException extends IcePatch2::FileAccessException

This exception is raised if an operation tries to use a file whose size is larger than 2.1 GB. Use the "large" versions of the operations instead.

Used By

- IcePatch2::FileServer::getFileCompressed
- IcePatch2::FileServer::getFileInfoSeq

Copyright 2016, ZeroC, Inc.

IcePatch2-LargeFileInfo

IcePatch2::LargeFileInfo

Overview

struct LargeFileInfo

Basic information about a single file.

Used By

- IcePatch2::LargeFileInfoSeq

Data Member Index

path — The pathname.
checksum — The SHA-1 checksum of the file.
size — The size of the compressed file in number of bytes.
executable — The executable flag.

Data Members

string path;

The pathname.

Ice::ByteSeq checksum;

The SHA-1 checksum of the file.

long size;

The size of the compressed file in number of bytes.

bool executable;

The executable flag.

Copyright 2016, ZeroC, Inc.

IcePatch2-PartitionOutOfRangeException

IcePatch2::PartitionOutOfRangeException

Overview

exception PartitionOutOfRangeException

A `partition` argument for was not in the range 0-255.

Used By

- IcePatch2::FileServer::getFileInfoSeq
- IcePatch2::FileServer::getLargeFileInfoSeq

Copyright 2016, ZeroC, Inc.

Ice Slice API

Ice

Overview

["objc:prefix:ICE"] module Ice

The Ice core library. Among many other features, the Ice core library manages all the communication tasks using an efficient protocol (including protocol compression and support for both TCP and UDP), provides a thread pool for multi-threaded servers, and additional functionality that supports high scalability.

Module Index

Instrumentation — The Instrumentation local interfaces enable observing a number of Ice core internal components (threads, connections, etc).

Class Index

ConnectionInfo — Base class providing access to the connection details.
EndpointInfo — Base class providing access to the endpoint details.
IPConnectionInfo — Provides access to the connection details of an IP connection
IPEndpointInfo — Provides access to the address details of a IP endpoint.
OpaqueEndpointInfo — Provides access to the details of an opaque endpoint.
TCPConnectionInfo — Provides access to the connection details of a TCP connection
TCPEndpointInfo — Provides access to a TCP endpoint information.
UDPConnectionInfo — Provides access to the connection details of a UDP connection
UDPEndpointInfo — Provides access to an UDP endpoint information.
WSConnectionInfo — Provides access to the connection details of a WebSocket connection
WSEndpointInfo — Provides access to a WebSocket endpoint information.

Interface Index

Communicator — The central object in Ice.
Connection — The user-level interface to a connection.
ConnectionCallback — An application can implement this interface to receive notifications when a connection closes or receives a heartbeat message.
Endpoint — The user-level interface to an endpoint.
ImplicitContext — An interface to associate implict contexts with communicators.
Locator — The Ice locator interface.
LocatorFinder — This inferface should be implemented by services implementing the Ice::Locator interface.
LocatorRegistry — The Ice locator registry interface.
Logger — The Ice message logger.
LoggerAdmin — The interface of the admin object that allows an Ice application the attach its RemoteLogger to the Logger of this admin object's Ice communicator.
ObjectAdapter — The object adapter provides an up-call interface from the Ice run time to the implementation of Ice objects.
ObjectFactory — A factory for objects.
Plugin — A communicator plug-in.
PluginManager — Each communicator has a plug-in manager to administer the set of plug-ins.
Process — An administrative interface for process management.
Properties — A property set used to configure Ice and Ice applications.
PropertiesAdmin — The PropertiesAdmin interface provides remote access to the properties of a communicator.
RemoteLogger — The Ice remote logger interface.
Router — The Ice router interface.
RouterFinder — This inferface should be implemented by services implementing the Ice::Router interface.
ServantLocator — A servant locator is called by an object adapter to locate a servant that is not found in its active servant map.

Exception Index

AdapterAlreadyActiveException — This exception is raised if a server tries to set endpoints for an adapter that is already active.

Copyright 2016, ZeroC, Inc.

AdapterNotFoundException — This exception is raised if an adapter cannot be found.
AlreadyRegisteredException — An attempt was made to register something more than once with the Ice run time.
BadMagicException — This exception indicates that a message did not start with the expected magic number ('I', 'c', 'e', 'P').
CloneNotImplementedException — An attempt was made to clone a class that does not support cloning.
CloseConnectionException — This exception indicates that the connection has been gracefully shut down by the server.
CloseTimeoutException — This exception indicates a connection closure timeout condition.
CollocationOptimizationException — This exception is raised if a feature is requested that is not supported with collocation optimization. *(Deprecated)*
CommunicatorDestroyedException — This exception is raised if the Communicator has been destroyed.
CompressionException — This exception indicates a problem with compressing or uncompressing data.
ConnectFailedException — This exception indicates connection failures.
ConnectTimeoutException — This exception indicates a connection establishment timeout condition.
ConnectionLostException — This exception indicates a lost connection.
ConnectionNotValidatedException — This exception is raised if a message is received over a connection that is not yet validated.
ConnectionRefusedException — This exception indicates a connection failure for which the server host actively refuses a connection.
ConnectionTimeoutException — This exception indicates that a connection has been shut down because it has been idle for some time.
DNSException — This exception indicates a DNS problem.
DatagramLimitException — A datagram exceeds the configured size.
EncapsulationException — This exception indicates a malformed data encapsulation.
EndpointParseException — This exception is raised if there was an error while parsing an endpoint.
EndpointSelectionTypeParseException — This exception is raised if there was an error while parsing an endpoint selection type.
FacetNotExistException — This exception is raised if no facet with the given name exists, but at least one facet with the given identity exists.
FeatureNotSupportedException — This exception is raised if an unsupported feature is used.
FileException — This exception indicates file errors.
FixedProxyException — This exception indicates that an attempt has been made to change the connection properties of a fixed proxy.
ForcedCloseConnectionException — This exception is raised by an operation call if the application forcefully closes the connection Connection::close.
IdentityParseException — This exception is raised if there was an error while parsing a stringified identity.
IllegalIdentityException — This exception is raised if an illegal identity is encountered.
IllegalMessageSizeException — This exception indicates that a message size is less than the minimum required size.
IllegalServantException — This exception is raised to reject an illegal servant (typically a null servant)
InitializationException — This exception is raised when a failure occurs during initialization.
InvalidReplicaGroupIdException — This exception is raised if the replica group provided by the server is invalid.
InvocationCanceledException — This exception indicates that an asynchronous invocation failed because it was canceled explicitly by the user using the `Ice::AsyncResult::cancel` method.
InvocationTimeoutException — This exception indicates that an invocation failed because it timed out.
MarshalException — This exception is raised for errors during marshaling or unmarshaling data.
MemoryLimitException — This exception is raised when Ice receives a request or reply message whose size exceeds the limit specified by the `Ice.MessageSizeMax` property.
NoEndpointException — This exception is raised if no suitable endpoint is available.
NoObjectFactoryException — This exception is raised if no suitable object factory was found during unmarshaling of a Slice class instance.
NotRegisteredException — An attempt was made to find or deregister something that is not registered with the Ice run time or Ice locator.
ObjectAdapterDeactivatedException — This exception is raised if an attempt is made to use a deactivated ObjectAdapter.
ObjectAdapterIdInUseException — This exception is raised if an ObjectAdapter cannot be activated.
ObjectNotExistException — This exception is raised if an object does not exist on the server, that is, if no facets with the given identity exist.
ObjectNotFoundException — This exception is raised if an object cannot be found.
OperationInterruptedException — This exception indicates a request was interrupted.
OperationNotExistException — This exception is raised if an operation for a given object does not exist on the server.
PluginInitializationException — This exception indicates that a failure occurred while initializing a plug-in.
ProtocolException — A generic exception base for all kinds of protocol error conditions.
ProxyParseException — This exception is raised if there was an error while parsing a stringified proxy.
ProxyUnmarshalException — This exception is raised if inconsistent data is received while unmarshaling a proxy.
RemoteLoggerAlreadyAttachedException — An exception thrown by LoggerAdmin::attachRemoteLogger to report that the provided RemoteLogger was previously attached to this LoggerAdmin.
RequestFailedException — This exception is raised if a request failed.
ResponseSentException — Indicates that the response to a request has already been sent; re-dispatching such a request is not possible.
SecurityException — This exception indicates a failure in a security subsystem, such as the IceSSL plug-in.
ServerNotFoundException — This exception is raised if a server cannot be found.
SocketException — This exception indicates socket errors.
StringConversionException — This exception is raised when a string conversion to or from UTF-8 fails during marshaling or unmarshaling.
SyscallException — This exception is raised if a system error occurred in the server or client process.
TimeoutException — This exception indicates a timeout condition.
TwowayOnlyException — The operation can only be invoked with a twoway request.
UnexpectedObjectException — This exception is raised if the type of an unmarshaled Slice class instance does not match its expected type.
UnknownException — This exception is raised if an operation call on a server raises an unknown exception.
UnknownLocalException — This exception is raised if an operation call on a server raises a local exception.
UnknownMessageException — This exception indicates that an unknown protocol message has been received.
UnknownReplyStatusException — This exception indicates that an unknown reply status has been received.

Copyright 2016, ZeroC, Inc.

UnknownRequestIdException — This exception indicates that a response for an unknown request ID has been received.
UnknownUserException — An operation raised an incorrect user exception.
UnmarshalOutOfBoundsException — This exception is raised if an out-of-bounds condition occurs during unmarshaling.
UnsupportedEncodingException — This exception indicates an unsupported data encoding version.
UnsupportedProtocolException — This exception indicates an unsupported protocol version.
VersionMismatchException — This exception is raised if the Ice library version does not match the version in the Ice header files.
VersionParseException — This exception is raised if there was an error while parsing a version.

Structure Index

ACM
Current — Information about the current method invocation for servers.
EncodingVersion — A version structure for the encoding version.
Identity — The identity of an Ice object.
LogMessage — A complete log message.
ProtocolVersion — A version structure for the protocol version.

Sequence Index

BoolSeq — A sequence of bools.
ByteSeq — A sequence of bytes.
DoubleSeq — A sequence of doubles.
EndpointSeq — A sequence of endpoints.
FloatSeq — A sequence of floats.
IdentitySeq — A sequence of identities.
IntSeq — A sequence of ints.
LogMessageSeq — A sequence of LogMessage.
LogMessageTypeSeq — A sequence of LogMessageType
LongSeq — A sequence of longs.
ObjectProxySeq — A sequence of object proxies.
ObjectSeq — A sequence of objects.
ShortSeq — A sequence of shorts.
StringSeq — A sequence of strings.

Dictionary Index

Context — A request context.
FacetMap — A mapping from facet name to servant.
HeaderDict
ObjectDict — A mapping between identities and Ice objects.
PropertyDict — A simple collection of properties, represented as a dictionary of key/value pairs.
SliceChecksumDict — A mapping from type IDs to Slice checksums.

Constant Index

SSLEndpointType — Uniquely identifies SSL endpoints.
TCPEndpointType — Uniquely identifies TCP endpoints.
UDPEndpointType — Uniquely identifies UDP endpoints.
WSEndpointType — Uniquely identifies TCP-based WebSocket endpoints.
WSSEndpointType — Uniquely identifies SSL-based WebSocket endpoints.

Enumeration Index

ACMClose
ACMHeartbeat
EndpointSelectionType — Determines the order in which the Ice run time uses the endpoints in a proxy when establishing a connection.
LogMessageType — An enumeration representing the different types of log messages.
OperationMode — The OperationMode determines the retry behavior an invocation in case of a (potentially) recoverable error.

Sequences

sequence<bool> BoolSeq

Copyright 2016, ZeroC, Inc.

A sequence of bools.

sequence<byte> ByteSeq

A sequence of bytes.

Used By

- BadMagicException::badMagic
- IcePatch2::ByteSeqSeq
- IcePatch2::FileInfo::checksum
- IcePatch2::FileServer::getChecksum
- IcePatch2::FileServer::getFileCompressed
- IcePatch2::FileServer::getLargeFileCompressed
- IcePatch2::LargeFileInfo::checksum
- OpaqueEndpointInfo::rawBytes

sequence<double> DoubleSeq

A sequence of doubles.

local sequence<Endpoint> EndpointSeq

A sequence of endpoints.

Used By

- ObjectAdapter::getEndpoints
- ObjectAdapter::getPublishedEndpoints

sequence<float> FloatSeq

A sequence of floats.

sequence<Identity> IdentitySeq

A sequence of identities.

Used By

- Glacier2::IdentitySet::add
- Glacier2::IdentitySet::get
- Glacier2::IdentitySet::remove
- IceStorm::Topic::getSubscribers

sequence<int> IntSeq

A sequence of ints.

["cpp:type:std::list<LogMessage>"] sequence<LogMessage> LogMessageSeq

A sequence of LogMessage.

Used By

- LoggerAdmin::getLog
- RemoteLogger::init

sequence<LogMessageType> LogMessageTypeSeq

A sequence of LogMessageType

Used By

Copyright 2016, ZeroC, Inc.

- LoggerAdmin::attachRemoteLogger
- LoggerAdmin::getLog

sequence<long> LongSeq

A sequence of longs.

sequence<Object*> ObjectProxySeq

A sequence of object proxies.

Used By

- IceGrid::Query::findAllObjectsByType
- IceGrid::Query::findAllReplicas
- IceGrid::TypeFilter::filter
- Router::addProxies

sequence<Object> ObjectSeq

A sequence of objects.

sequence<short> ShortSeq

A sequence of shorts.

sequence<string> StringSeq

A sequence of strings.

Used By

- Glacier2::SSLInfo::certs
- Glacier2::StringSet::add
- Glacier2::StringSet::get
- Glacier2::StringSet::remove
- IceBox::Service::start
- IceBox::ServiceObserver::servicesStarted
- IceBox::ServiceObserver::servicesStopped
- IceGrid::Admin::getAllAdapterIds
- IceGrid::Admin::getAllApplicationNames
- IceGrid::Admin::getAllNodeNames
- IceGrid::Admin::getAllRegistryNames
- IceGrid::Admin::getAllServerIds
- IceGrid::ApplicationUpdateDescriptor::removeNodes
- IceGrid::ApplicationUpdateDescriptor::removePropertySets
- IceGrid::ApplicationUpdateDescriptor::removeReplicaGroups
- IceGrid::ApplicationUpdateDescriptor::removeServerTemplates
- IceGrid::ApplicationUpdateDescriptor::removeServiceTemplates
- IceGrid::ApplicationUpdateDescriptor::removeVariables
- IceGrid::CommunicatorDescriptor::logs
- IceGrid::DistributionDescriptor::directories
- IceGrid::FileIterator::read
- IceGrid::NodeUpdateDescriptor::removePropertySets
- IceGrid::NodeUpdateDescriptor::removeServers
- IceGrid::NodeUpdateDescriptor::removeVariables
- IceGrid::PatchException::reasons
- IceGrid::PropertySetDescriptor::references
- IceGrid::ReplicaGroupFilter::filter
- IceGrid::ServerDescriptor::envs
- IceGrid::ServerDescriptor::options
- IceGrid::TemplateDescriptor::parameters
- IceMX::MetricsAdmin::getMetricsViewNames
- IceSSL::ConnectionInfo::certs
- LoggerAdmin::attachRemoteLogger

Copyright 2016, ZeroC, Inc.

- LoggerAdmin::getLog
- PluginManager::getPlugins
- Properties::getCommandLineOptions
- Properties::getPropertyAsListWithDefault
- Properties::getPropertyAsList
- Properties::parseCommandLineOptions
- Properties::parseIceCommandLineOptions

Dictionaries

dictionary<string, string> Context

A request context. `Context` is used to transmit metadata about a request from the server to the client, such as Quality-of-Service (QoS) parameters. Each operation on the client has a `Context` as its implicit final parameter.

Used By

- Current::ctx
- IceGrid::ReplicaGroupFilter::filter
- IceGrid::TypeFilter::filter
- ImplicitContext::getContext
- ImplicitContext::setContext
- Instrumentation::CommunicatorObserver::getInvocationObserver

local dictionary<string, Object> FacetMap

A mapping from facet name to servant.

Used By

- Communicator::findAllAdminFacets
- ObjectAdapter::findAllFacets
- ObjectAdapter::removeAllFacets

dictionary<string, string> HeaderDict

Used By

- IceSSL::WSSConnectionInfo::headers
- WSConnectionInfo::headers

local dictionary<**Identity**, Object> ObjectDict

A mapping between identities and Ice objects.

dictionary<string, string> PropertyDict

A simple collection of properties, represented as a dictionary of key/value pairs. Both key and value are strings.

Used By

- Communicator::proxyToProperty
- Properties::getPropertiesForPrefix
- PropertiesAdmin::getPropertiesForPrefix
- PropertiesAdmin::setProperties

See Also

- Properties::getPropertiesForPrefix

dictionary<string, string> SliceChecksumDict

A mapping from type IDs to Slice checksums. The dictionary allows verification at run time that client and server use matching Slice

Copyright 2016, ZeroC, Inc.

definitions.

Used By

- IceBox::ServiceManager::getSliceChecksums
- IceGrid::Admin::getSliceChecksums
- IceStorm::TopicManager::getSliceChecksums

Constants

const short SSLEndpointType = 2;

Uniquely identifies SSL endpoints.

const short TCPEndpointType = 1;

Uniquely identifies TCP endpoints.

const short UDPEndpointType = 3;

Uniquely identifies UDP endpoints.

const short WSEndpointType = 4;

Uniquely identifies TCP-based WebSocket endpoints.

const short WSSEndpointType = 5;

Uniquely identifies SSL-based WebSocket endpoints.

Copyright 2016, ZeroC, Inc.

Ice-ACM

Ice::ACM

Overview

local struct ACM

Used By

- Ice::Connection::getACM

Data Member Index

timeout
close
heartbeat

Data Members

int timeout;

Ice::ACMClose close;

Ice::ACMHeartbeat heartbeat;

Copyright 2016, ZeroC, Inc.

Ice-ACMClose

Ice::ACMClose

Overview

local enum ACMClose

Used By

- ACM::close
- Connection::setACM

Enumerator Index

CloseOff
CloseOnIdle
CloseOnInvocation
CloseOnInvocationAndIdle
CloseOnIdleForceful

Enumerators

CloseOff

CloseOnIdle

CloseOnInvocation

CloseOnInvocationAndIdle

CloseOnIdleForceful

Copyright 2016, ZeroC, Inc.

Ice-ACMHeartbeat

Ice::ACMHeartbeat

Overview

local enum ACMHeartbeat

Used By

- ACM::heartbeat
- Connection::setACM

Enumerator Index

HeartbeatOff
HeartbeatOnInvocation
HeartbeatOnIdle
HeartbeatAlways

Enumerators

HeartbeatOff

HeartbeatOnInvocation

HeartbeatOnIdle

HeartbeatAlways

Copyright 2016, ZeroC, Inc.

Ice-AdapterAlreadyActiveException

Ice::AdapterAlreadyActiveException

Overview

exception AdapterAlreadyActiveException

This exception is raised if a server tries to set endpoints for an adapter that is already active.

Used By

- Ice::LocatorRegistry::setAdapterDirectProxy
- Ice::LocatorRegistry::setReplicatedAdapterDirectProxy

Copyright 2016, ZeroC, Inc.

Ice-AdapterNotFoundException

Ice::AdapterNotFoundException

Overview

exception AdapterNotFoundException

This exception is raised if an adapter cannot be found.

Used By

- Ice::Locator::findAdapterById
- Ice::LocatorRegistry::setAdapterDirectProxy
- Ice::LocatorRegistry::setReplicatedAdapterDirectProxy

Copyright 2016, ZeroC, Inc.

Ice-AlreadyRegisteredException

Ice::AlreadyRegisteredException

Overview

local exception AlreadyRegisteredException

An attempt was made to register something more than once with the Ice run time. This exception is raised if an attempt is made to register a servant, servant locator, facet, object factory, plug-in, object adapter, object, or user exception factory more than once for the same ID.

Data Member Index

kindOfObject — The kind of object that is registered already: "servant", "servant locator", "object factory", "plug-in", "object adapter", "object", or "user exception factory".
id — The ID (or name) of the object that is registered already.

Data Members

string kindOfObject;

The kind of object that is registered already: "servant", "servant locator", "object factory", "plug-in", "object adapter", "object", or "user exception factory".

string id;

The ID (or name) of the object that is registered already.

Copyright 2016, ZeroC, Inc.

Ice-BadMagicException

Ice::BadMagicException

Overview

local exception BadMagicException extends Ice::ProtocolException

This exception indicates that a message did not start with the expected magic number ('I', 'c', 'e', 'P').

Data Member Index

badMagic — A sequence containing the first four bytes of the incorrect message.

Data Members

Ice::ByteSeq badMagic;

A sequence containing the first four bytes of the incorrect message.

Copyright 2016, ZeroC, Inc.

Ice-CloneNotImplementedException

Ice::CloneNotImplementedException

Overview

local exception CloneNotImplementedException

An attempt was made to clone a class that does not support cloning. This exception is raised if `ice_clone` is called on a class that is derived from an abstract Slice class (that is, a class containing operations), and the derived class does not provide an implementation of the `ice_clone` operation (C++ only).

Copyright 2016, ZeroC, Inc.

Ice-CloseConnectionException

Ice::CloseConnectionException

Overview

local exception CloseConnectionException extends Ice::ProtocolException

This exception indicates that the connection has been gracefully shut down by the server. The operation call that caused this exception has not been executed by the server. In most cases you will not get this exception, because the client will automatically retry the operation call in case the server shut down the connection. However, if upon retry the server shuts down the connection again, and the retry limit has been reached, then this exception is propagated to the application code.

Copyright 2016, ZeroC, Inc.

Ice-CloseTimeoutException

Ice::CloseTimeoutException

Overview

local exception CloseTimeoutException extends Ice::TimeoutException

This exception indicates a connection closure timeout condition.

Copyright 2016, ZeroC, Inc.

Ice-CollocationOptimizationException

Ice::CollocationOptimizationException

Overview

local exception CollocationOptimizationException

This exception is raised if a feature is requested that is not supported with collocation optimization.

this exception isn't used anymore by the Ice runtime

Copyright 2016, ZeroC, Inc.

Ice-Communicator

Ice::Communicator

Overview

["clr:implements:_System.IDisposable"] local interface Communicator

The central object in Ice. One or more communicators can be instantiated for an Ice application. Communicator instantiation is language-specific, and not specified in Slice code.

Used By

- Freeze::Connection::getCommunicator
- Ice::ObjectAdapter::getCommunicator
- IceBox::Service::start

See Also

- Ice::Logger
- Ice::ObjectAdapter
- Ice::Properties
- Ice::ObjectFactory

Operation Index

destroy — Destroy the communicator.
shutdown — Shuts down this communicator's server functionality, which includes the deactivation of all object adapters.
waitForShutdown — Wait until the application has called shutdown (or destroy).
isShutdown — Check whether communicator has been shut down.
stringToProxy — Convert a stringified proxy into a proxy.
proxyToString — Convert a proxy into a string.
propertyToProxy — Convert a set of proxy properties into a proxy.
proxyToProperty — Convert a proxy to a set of proxy properties.
stringToIdentity — Convert a string into an identity.
identityToString — Convert an identity into a string.
createObjectAdapter — Create a new object adapter.
createObjectAdapterWithEndpoints — Create a new object adapter with endpoints.
createObjectAdapterWithRouter — Create a new object adapter with a router.
addObjectFactory — Add an object factory to this communicator.
findObjectFactory — Find an object factory registered with this communicator.
getImplicitContext — Get the implicit context associated with this communicator.
getProperties — Get the properties for this communicator.
getLogger — Get the logger for this communicator.
getObserver — Get the observer resolver object for this communicator.
getDefaultRouter — Get the default router this communicator.
setDefaultRouter — Set a default router for this communicator.
getDefaultLocator — Get the default locator this communicator.
setDefaultLocator — Set a default Ice locator for this communicator.
getPluginManager — Get the plug-in manager for this communicator.
flushBatchRequests — Flush any pending batch requests for this communicator.
createAdmin — Add the Admin object with all its facets to the provided object adapter.
getAdmin — Get a proxy to the main facet of the Admin object.
addAdminFacet — Add a new facet to the Admin object.
removeAdminFacet — Remove the following facet to the Admin object.
findAdminFacet — Returns a facet of the Admin object.
findAllAdminFacets — Returns a map of all facets of the Admin object.

Operations

void destroy()

Destroy the communicator. This operation calls shutdown implicitly. Calling destroy cleans up memory, and shuts down this communicator's client functionality and destroys all object adapters. Subsequent calls to destroy are ignored.

Copyright 2016, ZeroC, Inc.

See Also

- shutdown
- Ice::ObjectAdapter::destroy

void shutdown()

Shuts down this communicator's server functionality, which includes the deactivation of all object adapters. (Attempts to use a deactivated object adapter raise Ice::ObjectAdapterDeactivatedException.) Subsequent calls to shutdown are ignored.

After shutdown returns, no new requests are processed. However, requests that have been started before shutdown was called might still be active. You can use waitForShutdown to wait for the completion of all requests.

See Also

- destroy
- waitForShutdown
- Ice::ObjectAdapter::deactivate

void waitForShutdown()

Wait until the application has called shutdown (or destroy). On the server side, this operation blocks the calling thread until all currently-executing operations have completed. On the client side, the operation simply block until another thread has called shutdown or destroy.

A typical use of this operation is to call it from the main thread, which then waits until some other thread calls shutdown. After shut-down is complete, the main thread returns and can do some cleanup work before it finally calls destroy to shut down the client functionality, and then exits the application.

See Also

- shutdown
- destroy
- Ice::ObjectAdapter::waitForDeactivate

bool isShutdown()

Check whether communicator has been shut down.

Return Value

True if the communicator has been shut down; false otherwise.

See Also

- shutdown

Object* stringToProxy(string str)

Convert a stringified proxy into a proxy. For example, `MyCategory/MyObject:tcp -h some_host -p 10000` creates a proxy that refers to the Ice object having an identity with a name "MyObject" and a category "MyCategory", with the server running on host "some_host", port 10000. If the stringified proxy does not parse correctly, the operation throws one of Ice::ProxyParseException, Ice::EndpointParseException, or Ice::IdentityParseException. An appendix in the Ice manual provides a detailed description of the syntax supported by stringified proxies.

Parameters

`str` — The stringified proxy to convert into a proxy.

Return Value

The proxy, or nil if `str` is an empty string.

See Also

- proxyToString

Copyright 2016, ZeroC, Inc.

string proxyToString(Object* obj)

Convert a proxy into a string.

Parameters

`obj` — The proxy to convert into a stringified proxy.

Return Value

The stringified proxy, or an empty string if `obj` is nil.

See Also

- stringToProxy

Object* propertyToProxy(string property)

Convert a set of proxy properties into a proxy. The "base" name supplied in the `property` argument refers to a property containing a stringified proxy, such as `MyProxy=id:tcp -h localhost -p 10000`. Additional properties configure local settings for the proxy, such as `MyProxy.PreferSecure=1`. The "Properties" appendix in the Ice manual describes each of the supported proxy properties.

Parameters

`property` — The base property name.

Return Value

The proxy.

Ice::PropertyDict proxyToProperty(Object* proxy, string property)

Convert a proxy to a set of proxy properties.

Parameters

`proxy` — The proxy.
`property` — The base property name.

Return Value

The property set.

Ice::Identity stringToIdentity(string str)

Convert a string into an identity. If the string does not parse correctly, the operation throws Ice::IdentityParseException.

Parameters

`str` — The string to convert into an identity.

Return Value

The identity.

See Also

- identityToString

string identityToString(Ice::Identity ident)

Convert an identity into a string.

Parameters

Copyright 2016, ZeroC, Inc.

`ident` — The identity to convert into a string.

Return Value

The "stringified" identity.

See Also

- stringToIdentity

Ice::ObjectAdapter createObjectAdapter(string name)

Create a new object adapter. The endpoints for the object adapter are taken from the property *`name`*`.Endpoints`.

It is legal to create an object adapter with the empty string as its name. Such an object adapter is accessible via bidirectional connections or by collocated invocations that originate from the same communicator as is used by the adapter.

Attempts to create a named object adapter for which no configuration can be found raise Ice::InitializationException.

Parameters

`name` — The object adapter name.

Return Value

The new object adapter.

See Also

- createObjectAdapterWithEndpoints
- Ice::ObjectAdapter
- Ice::Properties

Ice::ObjectAdapter createObjectAdapterWithEndpoints(string name, string endpoints)

Create a new object adapter with endpoints. This operation sets the property *`name`*`.Endpoints`, and then calls createObjectAdapter. It is provided as a convenience function.

Calling this operation with an empty name will result in a UUID being generated for the name.

Parameters

`name` — The object adapter name.
`endpoints` — The endpoints for the object adapter.

Return Value

The new object adapter.

See Also

- createObjectAdapter
- Ice::ObjectAdapter
- Ice::Properties

Ice::ObjectAdapter createObjectAdapterWithRouter(string name, Ice::Router* rtr)

Create a new object adapter with a router. This operation creates a routed object adapter.

Calling this operation with an empty name will result in a UUID being generated for the name.

Parameters

`name` — The object adapter name.
`rtr` — The router.

Copyright 2016, ZeroC, Inc.

Return Value

The new object adapter.

See Also

- createObjectAdapter
- Ice::ObjectAdapter
- Ice::Properties

void addObjectFactory(Ice::ObjectFactory factory, string id)

Add an object factory to this communicator. Installing a factory with an id for which a factory is already registered throws Ice::AlreadyRegisteredException.

When unmarshaling an Ice object, the Ice run time reads the most-derived type id off the wire and attempts to create an instance of the type using a factory. If no instance is created, either because no factory was found, or because all factories returned nil, the behavior of the Ice run time depends on the format with which the object was marshaled:

If the object uses the "sliced" format, Ice ascends the class hierarchy until it finds a type that is recognized by a factory, or it reaches the least-derived type. If no factory is found that can create an instance, the run time throws Ice::NoObjectFactoryException.

If the object uses the "compact" format, Ice immediately raises Ice::NoObjectFactoryException.

The following order is used to locate a factory for a type:

1. The Ice run-time looks for a factory registered specifically for the type.
2. If no instance has been created, the Ice run-time looks for the default factory, which is registered with an empty type id.
3. If no instance has been created by any of the preceding steps, the Ice run-time looks for a factory that may have been statically generated by the language mapping for non-abstract classes.

Parameters

`factory` — The factory to add.
`id` — The type id for which the factory can create instances, or an empty string for the default factory.

See Also

- findObjectFactory
- Ice::ObjectFactory

Ice::ObjectFactory findObjectFactory(string id)

Find an object factory registered with this communicator.

Parameters

`id` — The type id for which the factory can create instances, or an empty string for the default factory.

Return Value

The object factory, or null if no object factory was found for the given id.

See Also

- addObjectFactory
- Ice::ObjectFactory

Ice::ImplicitContext getImplicitContext()

Get the implicit context associated with this communicator.

Return Value

The implicit context associated with this communicator; returns null when the property Ice.ImplicitContext is not set or is set to None.

Ice::Properties getProperties()

Copyright 2016, ZeroC, Inc.

Get the properties for this communicator.

Return Value

This communicator's properties.

See Also

- Ice::Properties

Ice::Logger getLogger()

Get the logger for this communicator.

Return Value

This communicator's logger.

See Also

- Ice::Logger

Ice::Instrumentation::CommunicatorObserver getObserver()

Get the observer resolver object for this communicator.

Return Value

This communicator's observer resolver object.

Ice::Router* getDefaultRouter()

Get the default router this communicator.

Return Value

The default router for this communicator.

See Also

- setDefaultRouter
- Ice::Router

void setDefaultRouter(Ice::Router* rtr)

Set a default router for this communicator. All newly created proxies will use this default router. To disable the default router, null can be used. Note that this operation has no effect on existing proxies.

You can also set a router for an individual proxy by calling the operation `ice_router` on the proxy.

Parameters

`rtr` — The default router to use for this communicator.

See Also

- getDefaultRouter
- createObjectAdapterWithRouter
- Ice::Router

Ice::Locator* getDefaultLocator()

Get the default locator this communicator.

Return Value

Copyright 2016, ZeroC, Inc.

The default locator for this communicator.

See Also

- setDefaultLocator
- Ice::Locator

void setDefaultLocator(Ice::Locator* loc)

Set a default Ice locator for this communicator. All newly created proxy and object adapters will use this default locator. To disable the default locator, null can be used. Note that this operation has no effect on existing proxies or object adapters.

You can also set a locator for an individual proxy by calling the operation `ice_locator` on the proxy, or for an object adapter by calling the operation Ice::ObjectAdapter::setLocator on the object adapter.

Parameters

`loc` — The default locator to use for this communicator.

See Also

- getDefaultLocator
- Ice::Locator
- Ice::ObjectAdapter::setLocator

Ice::PluginManager getPluginManager()

Get the plug-in manager for this communicator.

Return Value

This communicator's plug-in manager.

See Also

- Ice::PluginManager

["async"] void flushBatchRequests()

Flush any pending batch requests for this communicator. This means all batch requests invoked on fixed proxies for all connections associated with the communicator. Any errors that occur while flushing a connection are ignored.

Object* createAdmin(Ice::ObjectAdapter adminAdapter, Ice::Identity adminId)

Add the Admin object with all its facets to the provided object adapter. If Ice.Admin.ServerId is set and the provided object adapter has a Ice::Locator, createAdmin registers the Admin's Process facet with the Ice::Locator's Ice::LocatorRegistry.

createAdmin call only be called once; subsequent calls raise Ice::InitializationException.

Parameters

`adminAdapter` — The object adapter used to host the Admin object; if null and Ice.Admin.Endpoints is set, create, activate and use the Ice.Admin object adapter.
`adminId` — The identity of the Admin object.

Return Value

A proxy to the main ("") facet of the Admin object. Never returns a null proxy.

See Also

- getAdmin
- Ice::LocatorRegistry::setServerProcessProxy

Object* getAdmin()

Copyright 2016, ZeroC, Inc.

Get a proxy to the main facet of the Admin object. getAdmin also creates the Admin object and creates and activates the Ice.Admin object adapter to host this Admin object if Ice.Admin.Enpoints is set. The identity of the Admin object created by getAdmin is **`UNRECOGNIZED MARKUP: value`**/admin, or **`UNRECOGNIZED MARKUP: UUID`**/admin when Ice.Admin.InstanceName is not set.

If Ice.Admin.DelayCreation is 0 or not set, getAdmin is called by the communicator initialization, after initialization of all plugins.

Return Value

A proxy to the main ("") facet of the Admin object, or a null proxy if no Admin object is configured.

See Also

- createAdmin

void addAdminFacet(Object servant, string facet)

Add a new facet to the Admin object. Adding a servant with a facet that is already registered throws Ice::AlreadyRegisteredException.

Parameters

`servant` — The servant that implements the new Admin facet.
`facet` — The name of the new Admin facet.

Object removeAdminFacet(string facet)

Remove the following facet to the Admin object. Removing a facet that was not previously registered throws Ice::NotRegisteredException.

Parameters

`facet` — The name of the Admin facet.

Return Value

The servant associated with this Admin facet.

Object findAdminFacet(string facet)

Returns a facet of the Admin object.

Parameters

`facet` — The name of the Admin facet.

Return Value

The servant associated with this Admin facet, or null if no facet is registered with the given name.

Ice::FacetMap findAllAdminFacets()

Returns a map of all facets of the Admin object.

Return Value

A collection containing all the facet names and servants of the Admin object.

See Also

- findAdminFacet

Copyright 2016, ZeroC, Inc.

Ice-CommunicatorDestroyedException

Ice::CommunicatorDestroyedException

Overview

local exception CommunicatorDestroyedException

This exception is raised if the Ice::Communicator has been destroyed.

See Also

- Ice::Communicator::destroy

Copyright 2016, ZeroC, Inc.

Ice-CompressionException

Ice::CompressionException

Overview

local exception CompressionException extends Ice::ProtocolException

This exception indicates a problem with compressing or uncompressing data.

Copyright 2016, ZeroC, Inc.

Ice-ConnectFailedException

Ice::ConnectFailedException

Overview

local exception ConnectFailedException extends Ice::SocketException

This exception indicates connection failures.

Derived Exceptions

- Ice::ConnectionRefusedException

Copyright 2016, ZeroC, Inc.

Ice-Connection

Ice::Connection

Overview

local interface Connection

The user-level interface to a connection.

Used By

- Ice::ConnectionCallback::closed
- Ice::ConnectionCallback::heartbeat
- Ice::Current::con
- IceGrid::ReplicaGroupFilter::filter
- IceGrid::TypeFilter::filter

Operation Index

close — Close a connection, either gracefully or forcefully.
createProxy — Create a special proxy that always uses this connection.
setAdapter — Explicitly set an object adapter that dispatches requests that are received over this connection.
getAdapter — Get the object adapter that dispatches requests for this connection.
getEndpoint — Get the endpoint from which the connection was created.
flushBatchRequests — Flush any pending batch requests for this connection.
setCallback — Set callback on the connection.
setACM — Set the active connection management parameters.
getACM — Get the ACM parameters.
type — Return the connection type.
timeout — Get the timeout for the connection.
toString — Return a description of the connection as human readable text, suitable for logging or error messages.
getInfo — Returns the connection information.
setBufferSize — Set the connectiion buffer receive/send size.

Operations

void close(bool force)

Close a connection, either gracefully or forcefully. If a connection is closed forcefully, it closes immediately, without sending the relevant close connection protocol messages to the peer and waiting for the peer to acknowledge these protocol messages.

Parameters

`force` — If true, close forcefully. Otherwise the connection is closed gracefully.

Object* createProxy(**Ice::Identity** id)

Create a special proxy that always uses this connection. This can be used for callbacks from a server to a client if the server cannot directly establish a connection to the client, for example because of firewalls. In this case, the server would create a proxy using an already established connection from the client.

Parameters

`id` — The identity for which a proxy is to be created.

Return Value

A proxy that matches the given identity and uses this connection.

See Also

- setAdapter

Copyright 2016, ZeroC, Inc.

void setAdapter(Ice::ObjectAdapter adapter)

Explicitly set an object adapter that dispatches requests that are received over this connection. A client can invoke an operation on a server using a proxy, and then set an object adapter for the outgoing connection that is used by the proxy in order to receive callbacks. This is useful if the server cannot establish a connection back to the client, for example because of firewalls.

Parameters

`adapter` — The object adapter that should be used by this connection to dispatch requests. The object adapter must be activated. When the object adapter is deactivated, it is automatically removed from the connection.

See Also

- createProxy
- setAdapter

Ice::ObjectAdapter getAdapter()

Get the object adapter that dispatches requests for this connection.

Return Value

The object adapter that dispatches requests for the connection, or null if no adapter is set.

See Also

- setAdapter

Ice::Endpoint getEndpoint()

Get the endpoint from which the connection was created.

Return Value

The endpoint from which the connection was created.

["async"] void flushBatchRequests()

Flush any pending batch requests for this connection. This means all batch requests invoked on fixed proxies associated with the connection.

void setCallback(Ice::ConnectionCallback callback)

Set callback on the connection. The callback is called by the connection when it's closed. The callback is called from the Ice thread pool associated with the connection.

Parameters

`callback` — The connection callback object.

["java:optional"] void setACM(int timeout, Ice::ACMClose close, Ice::ACMHeartbeat heartbeat)

Set the active connection management parameters.

Parameters

`timeout` — The timeout value in milliseconds.
`close` — The close condition
`heartbeat` — The hertbeat condition

Ice::ACM getACM()

Get the ACM parameters.

Copyright 2016, ZeroC, Inc.

Return Value

The ACM parameters.

string type()

Return the connection type. This corresponds to the endpoint type, i.e., "tcp", "udp", etc.

Return Value

The type of the connection.

int timeout()

Get the timeout for the connection.

Return Value

The connection's timeout.

string toString()

Return a description of the connection as human readable text, suitable for logging or error messages.

Return Value

The description of the connection as human readable text.

Ice::ConnectionInfo getInfo()

Returns the connection information.

Return Value

The connection information.

void setBufferSize(int rcvSize, int sndSize)

Set the connectiion buffer receive/send size.

Parameters

`rcvSize` — The connection receive buffer size.
`sndSize` — The connection send buffer size.

Copyright 2016, ZeroC, Inc.

Ice-ConnectionCallback

Ice::ConnectionCallback

Overview

local interface ConnectionCallback

An application can implement this interface to receive notifications when a connection closes or receives a heartbeat message.

Used By

- Ice::Connection::setCallback

See Also

- Ice::Connection::setCallback

Operation Index

heartbeat — This method is called by the the connection when a heartbeat is received from the peer.
closed — This method is called by the the connection when the connection is closed.

Operations

void heartbeat(**Ice::Connection** con)

This method is called by the the connection when a heartbeat is received from the peer.

void closed(**Ice::Connection** con)

This method is called by the the connection when the connection is closed.

Copyright 2016, ZeroC, Inc.

Ice-ConnectionInfo

Ice::ConnectionInfo

Overview

local class ConnectionInfo

Base class providing access to the connection details. *

Derived Classes and Interfaces

- Ice::IPConnectionInfo

Used By

- Ice::Connection::getInfo
- Ice::Instrumentation::CommunicatorObserver::getConnectionObserver
- Ice::Instrumentation::InvocationObserver::getRemoteObserver

Data Member Index

incoming — Whether or not the connection is an incoming or outgoing connection.
adapterName — The name of the adapter associated with the connection.
connectionId — The connection id.
rcvSize — The connection buffer receive size.
sndSize — The connection buffer send size.

Data Members

bool incoming;

Whether or not the connection is an incoming or outgoing connection.

string adapterName;

The name of the adapter associated with the connection.

string connectionId;

The connection id.

int rcvSize;

The connection buffer receive size.

int sndSize;

The connection buffer send size.

Copyright 2016, ZeroC, Inc.

Ice-ConnectionLostException

Ice::ConnectionLostException

Overview

local exception ConnectionLostException extends Ice::SocketException

This exception indicates a lost connection.

Copyright 2016, ZeroC, Inc.

Ice-ConnectionNotValidatedException

Ice::ConnectionNotValidatedException

Overview

local exception ConnectionNotValidatedException extends Ice::ProtocolException

This exception is raised if a message is received over a connection that is not yet validated.

Copyright 2016, ZeroC, Inc.

Ice-ConnectionRefusedException

Ice::ConnectionRefusedException

Overview

local exception ConnectionRefusedException extends Ice::ConnectFailedException

This exception indicates a connection failure for which the server host actively refuses a connection.

Copyright 2016, ZeroC, Inc.

Ice-ConnectionTimeoutException

Ice::ConnectionTimeoutException

Overview

local exception ConnectionTimeoutException extends Ice::TimeoutException

This exception indicates that a connection has been shut down because it has been idle for some time.

Copyright 2016, ZeroC, Inc.

Ice-ConnectTimeoutException

Ice::ConnectTimeoutException

Overview

local exception ConnectTimeoutException extends Ice::TimeoutException

This exception indicates a connection establishment timeout condition.

Copyright 2016, ZeroC, Inc.

Ice-Current

Ice::Current

Overview

local struct Current

Information about the current method invocation for servers. Each operation on the server has a `Current` as its implicit final parameter. `Current` is mostly used for Ice services. Most applications ignore this parameter.

Used By

- Ice::Instrumentation::CommunicatorObserver::getDispatchObserver
- Ice::ServantLocator::finished
- Ice::ServantLocator::locate

Data Member Index

adapter — The object adapter.
con — Information about the connection over which the current method invocation was received.
id — The Ice object identity.
facet — The facet.
operation — The operation name.
mode — The mode of the operation.
ctx — The request context, as received from the client.
requestId — The request id unless oneway (0) or collocated (-1).
encoding — The encoding version used to encode the input and output parameters.

Data Members

Ice::ObjectAdapter adapter;

The object adapter.

Ice::Connection con;

Information about the connection over which the current method invocation was received. If the invocation is direct due to collocation optimization, this value is set to null.

Ice::Identity id;

The Ice object identity.

string facet;

The facet.

string operation;

The operation name.

Ice::OperationMode mode;

The mode of the operation.

Ice::Context ctx;

The request context, as received from the client.

int requestId;

Copyright 2016, ZeroC, Inc.

The request id unless oneway (0) or collocated (-1).

Ice::EncodingVersion encoding;

The encoding version used to encode the input and output parameters.

Copyright 2016, ZeroC, Inc.

Ice-DatagramLimitException

Ice::DatagramLimitException

Overview

local exception DatagramLimitException extends Ice::ProtocolException

A datagram exceeds the configured size. This exception is raised if a datagram exceeds the configured send or receive buffer size, or exceeds the maximum payload size of a UDP packet (65507 bytes).

Copyright 2016, ZeroC, Inc.

Ice-DNSException

Ice::DNSException

Overview

local exception DNSException

This exception indicates a DNS problem. For details on the cause, error should be inspected.

Data Member Index

error — The error number describing the DNS problem.
host — The host name that could not be resolved.

Data Members

int error;

The error number describing the DNS problem. For C++ and Unix, this is equivalent to `h_errno`. For C++ and Windows, this is the value returned by `WSAGetLastError()`.

string host;

The host name that could not be resolved.

Copyright 2016, ZeroC, Inc.

Ice-EncapsulationException

Ice::EncapsulationException

Overview

local exception EncapsulationException extends Ice::MarshalException

This exception indicates a malformed data encapsulation.

Copyright 2016, ZeroC, Inc.

Ice-EncodingVersion

Ice::EncodingVersion

Overview

struct EncodingVersion

A version structure for the encoding version.

Used By

- Freeze::Connection::getEncoding
- Ice::Current::encoding
- Ice::OpaqueEndpointInfo::rawEncoding
- Ice::UnsupportedEncodingException::bad
- Ice::UnsupportedEncodingException::supported

Data Member Index

major
minor

Data Members

byte major;

byte minor;

Copyright 2016, ZeroC, Inc.

Ice-Endpoint

Ice::Endpoint

Overview

local interface Endpoint

The user-level interface to an endpoint.

Used By

- Ice::Connection::getEndpoint
- Ice::EndpointSeq
- Ice::Instrumentation::CommunicatorObserver::getConnectionEstablishmentObserver
- Ice::Instrumentation::CommunicatorObserver::getConnectionObserver
- Ice::Instrumentation::CommunicatorObserver::getEndpointLookupObserver
- Ice::Instrumentation::InvocationObserver::getRemoteObserver

Operation Index

toString — Return a string representation of the endpoint.
getInfo — Returns the endpoint information.

Operations

string toString()

Return a string representation of the endpoint.

Return Value

The string representation of the endpoint.

Ice::EndpointInfo getInfo()

Returns the endpoint information.

Return Value

The endpoint information class.

Copyright 2016, ZeroC, Inc.

Ice-EndpointInfo

Ice::EndpointInfo

Overview

local class EndpointInfo

Base class providing access to the endpoint details.

Derived Classes and Interfaces

- Ice::IPEndpointInfo
- Ice::OpaqueEndpointInfo

Used By

- Ice::Endpoint::getInfo

Operation Index

type — Returns the type of the endpoint.
datagram — Returns true if this endpoint is a datagram endpoint.
secure — Returns true if this endpoint is a secure endpoint.

Data Member Index

timeout — The timeout for the endpoint in milliseconds.
compress — Specifies whether or not compression should be used if available when using this endpoint.

Operations

short type()

Returns the type of the endpoint.

bool datagram()

Returns true if this endpoint is a datagram endpoint.

bool secure()

Returns true if this endpoint is a secure endpoint.

Data Members

int timeout;

The timeout for the endpoint in milliseconds. 0 means non-blocking, -1 means no timeout.

bool compress;

Specifies whether or not compression should be used if available when using this endpoint.

Copyright 2016, ZeroC, Inc.

Ice-EndpointParseException

Ice::EndpointParseException

Overview

local exception EndpointParseException

This exception is raised if there was an error while parsing an endpoint.

Data Member Index

str — Describes the failure and includes the string that could not be parsed.

Data Members

string str;

Describes the failure and includes the string that could not be parsed.

Copyright 2016, ZeroC, Inc.

Ice-EndpointSelectionType

Ice::EndpointSelectionType

Overview

enum EndpointSelectionType

Determines the order in which the Ice run time uses the endpoints in a proxy when establishing a connection.

Enumerator Index

Random — `Random` causes the endpoints to be arranged in a random order.
Ordered — `Ordered` forces the Ice run time to use the endpoints in the order they appeared in the proxy.

Enumerators

Random

`Random` causes the endpoints to be arranged in a random order.

Ordered

`Ordered` forces the Ice run time to use the endpoints in the order they appeared in the proxy.

Copyright 2016, ZeroC, Inc.

Ice-EndpointSelectionTypeParseException

Ice::EndpointSelectionTypeParseException

Overview

local exception EndpointSelectionTypeParseException

This exception is raised if there was an error while parsing an endpoint selection type.

Data Member Index

str — Describes the failure and includes the string that could not be parsed.

Data Members

string str;

Describes the failure and includes the string that could not be parsed.

Copyright 2016, ZeroC, Inc.

Ice-FacetNotExistException

Ice::FacetNotExistException

Overview

local exception FacetNotExistException extends Ice::RequestFailedException

This exception is raised if no facet with the given name exists, but at least one facet with the given identity exists.

Copyright 2016, ZeroC, Inc.

Ice-FeatureNotSupportedException

Ice::FeatureNotSupportedException

Overview

local exception FeatureNotSupportedException

This exception is raised if an unsupported feature is used. The unsupported feature string contains the name of the unsupported feature

Data Member Index

unsupportedFeature — The name of the unsupported feature.

Data Members

string unsupportedFeature;

The name of the unsupported feature.

Copyright 2016, ZeroC, Inc.

Ice-FileException

Ice::FileException

Overview

local exception FileException extends Ice::SyscallException

This exception indicates file errors.

Data Member Index

path — The path of the file responsible for the error.

Data Members

string path;

The path of the file responsible for the error.

Copyright 2016, ZeroC, Inc.

Ice-FixedProxyException

Ice::FixedProxyException

Overview

local exception FixedProxyException

This exception indicates that an attempt has been made to change the connection properties of a fixed proxy.

Copyright 2016, ZeroC, Inc.

Ice-ForcedCloseConnectionException

Ice::ForcedCloseConnectionException

Overview

local exception ForcedCloseConnectionException extends Ice::ProtocolException

This exception is raised by an operation call if the application forcefully closes the connection Ice::Connection::close.

See Also

- Ice::Connection::close

Copyright 2016, ZeroC, Inc.

Ice-Identity

Ice::Identity

Overview

struct Identity

The identity of an Ice object. In a proxy, an empty name denotes a nil proxy. An identity with an empty name and a non-empty category is illegal. You cannot add a servant with an empty name to the Active Servant Map.

Used By

- Freeze::BackgroundSaveEvictor::keepFacet
- Freeze::BackgroundSaveEvictor::keep
- Freeze::BackgroundSaveEvictor::releaseFacet
- Freeze::BackgroundSaveEvictor::release
- Freeze::Evictor::addFacet
- Freeze::Evictor::add
- Freeze::Evictor::hasFacet
- Freeze::Evictor::hasObject
- Freeze::Evictor::removeFacet
- Freeze::Evictor::remove
- Freeze::EvictorIterator::next
- Freeze::ServantInitializer::initialize
- Ice::Communicator::createAdmin
- Ice::Communicator::identityToString
- Ice::Communicator::stringToIdentity
- Ice::Connection::createProxy
- Ice::Current::id
- Ice::IdentitySeq
- Ice::IllegalIdentityException::id
- Ice::Locator::findObjectById
- Ice::ObjectAdapter::addFacet
- Ice::ObjectAdapter::add
- Ice::ObjectAdapter::createDirectProxy
- Ice::ObjectAdapter::createIndirectProxy
- Ice::ObjectAdapter::createProxy
- Ice::ObjectAdapter::findAllFacets
- Ice::ObjectAdapter::findFacet
- Ice::ObjectAdapter::find
- Ice::ObjectAdapter::removeAllFacets
- Ice::ObjectAdapter::removeFacet
- Ice::ObjectAdapter::remove
- Ice::ObjectDict
- Ice::RequestFailedException::id
- IceGrid::Admin::getObjectInfo
- IceGrid::Admin::removeObject
- IceGrid::AdminSession::setObserversByIdentity
- IceGrid::ObjectDescriptor::id
- IceGrid::ObjectExistsException::id
- IceGrid::ObjectNotRegisteredException::id
- IceGrid::ObjectObserver::objectRemoved
- IceGrid::ObserverAlreadyRegisteredException::id
- IceGrid::Query::findObjectById
- IceGrid::RegistryPluginFacade::getObjectInfo
- IceGrid::Session::allocateObjectById
- IceGrid::Session::releaseObject

See Also

- Ice::ServantLocator
- Ice::ObjectAdapter::addServantLocator

Copyright 2016, ZeroC, Inc.

Data Member Index

name — The name of the Ice object.
category — The Ice object category.

Data Members

string name;

The name of the Ice object.

string category;

The Ice object category.

Copyright 2016, ZeroC, Inc.

Ice-IdentityParseException

Ice::IdentityParseException

Overview

local exception IdentityParseException

This exception is raised if there was an error while parsing a stringified identity.

Data Member Index

str — Describes the failure and includes the string that could not be parsed.

Data Members

string str;

Describes the failure and includes the string that could not be parsed.

Copyright 2016, ZeroC, Inc.

Ice-IllegalIdentityException

Ice::IllegalIdentityException

Overview

local exception IllegalIdentityException

This exception is raised if an illegal identity is encountered.

Data Member Index

id — The illegal identity.

Data Members

Ice::Identity id;

The illegal identity.

Copyright 2016, ZeroC, Inc.

Ice-IllegalMessageSizeException

Ice::IllegalMessageSizeException

Overview

local exception IllegalMessageSizeException extends Ice::ProtocolException

This exception indicates that a message size is less than the minimum required size.

Copyright 2016, ZeroC, Inc.

Ice-IllegalServantException

Ice::IllegalServantException

Overview

local exception IllegalServantException

This exception is raised to reject an illegal servant (typically a null servant)

Data Member Index

reason — Describes why this servant is illegal.

Data Members

string reason;

Describes why this servant is illegal.

Copyright 2016, ZeroC, Inc.

Ice-ImplicitContext

Ice::ImplicitContext

Overview

local interface ImplicitContext

An interface to associate implict contexts with communicators. When you make a remote invocation without an explicit context parameter, Ice uses the per-proxy context (if any) combined with the `ImplicitContext` associated with the communicator.

Ice provides several implementations of `ImplicitContext`. The implementation used depends on the value of the `Ice.ImplicitContext` property.
`None` (default) — No implicit context at all.
`PerThread` — The implementation maintains a context per thread.
`Shared` — The implementation maintains a single context shared by all threads.

`ImplicitContext` also provides a number of operations to create, update or retrieve an entry in the underlying context without first retrieving a copy of the entire context. These operations correspond to a subset of the `java.util.Map` methods, with `java.lang.Object` replaced by `string` and null replaced by the empty-string.

Used By

- Ice::Communicator::getImplicitContext

Operation Index

getContext — Get a copy of the underlying context.
setContext — Set the underlying context.
containsKey — Check if this key has an associated value in the underlying context.
get — Get the value associated with the given key in the underlying context.
put — Create or update a key/value entry in the underlying context.
remove — Remove the entry for the given key in the underlying context.

Operations

Ice::Context getContext()

Get a copy of the underlying context.

Return Value

A copy of the underlying context.

void setContext(Ice::Context newContext)

Set the underlying context.

Parameters

`newContext` — The new context.

bool containsKey(string key)

Check if this key has an associated value in the underlying context.

Parameters

`key` — The key.

Return Value

Copyright 2016, ZeroC, Inc.

True if the key has an associated value, False otherwise.

string get(string key)

Get the value associated with the given key in the underlying context. Returns an empty string if no value is associated with the key. containsKey allows you to distinguish between an empty-string value and no value at all.

Parameters

`key` — The key.

Return Value

The value associated with the key.

string put(string key, string value)

Create or update a key/value entry in the underlying context.

Parameters

`key` — The key.
`value` — The value.

Return Value

The previous value associated with the key, if any.

string remove(string key)

Remove the entry for the given key in the underlying context.

Parameters

`key` — The key.

Return Value

The value associated with the key, if any.

Copyright 2016, ZeroC, Inc.

Ice-InitializationException

Ice::InitializationException

Overview

local exception InitializationException

This exception is raised when a failure occurs during initialization.

Data Member Index

reason — The reason for the failure.

Data Members

string reason;

The reason for the failure.

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation

Ice::Instrumentation

Overview

["objc:prefix:ICEINSTRUMENTATION"] module Instrumentation

The Instrumentation local interfaces enable observing a number of Ice core internal components (threads, connections, etc).

Interface Index

ChildInvocationObserver — The child invocation observer to instrument remote or collocated invocations.
CollocatedObserver — The collocated observer to instrument invocations that are collocated.
CommunicatorObserver — The communicator observer interface used by the Ice run-time to obtain and update observers for its observable objects.
ConnectionObserver — The connection observer interface to instrument Ice connections.
DispatchObserver — The dispatch observer to instrument servant dispatch.
InvocationObserver — The invocation observer to instrument invocations on proxies.
Observer — The object observer interface used by instrumented objects to notify the observer of their existence.
ObserverUpdater — The observer updater interface.
RemoteObserver — The remote observer to instrument invocations that are sent over the wire.
ThreadObserver — The thread observer interface to instrument Ice threads.

Enumeration Index

ConnectionState — The state of an Ice connection.
ThreadState — The thread state enumeration keeps track of the different possible states of Ice threads.

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation-ChildInvocationObserver

Ice::Instrumentation::ChildInvocationObserver

Overview

local interface ChildInvocationObserver extends Ice::Instrumentation::Observer

The child invocation observer to instrument remote or collocated invocations.
Derived Classes and Interfaces

- Ice::Instrumentation::CollocatedObserver
- Ice::Instrumentation::RemoteObserver

Operation Index

reply — Reply notification.

Operations

void reply(int size)

Reply notification.
Parameters

`size` — The size of the reply.

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation-CollocatedObserver

Ice::Instrumentation::CollocatedObserver

Overview

local interface CollocatedObserver extends Ice::Instrumentation::ChildInvocationObserver

The collocated observer to instrument invocations that are collocated.
Used By

- Ice::Instrumentation::InvocationObserver::getCollocatedObserver

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation-CommunicatorObserver

Ice::Instrumentation::CommunicatorObserver

Overview

local interface CommunicatorObserver

The communicator observer interface used by the Ice run-time to obtain and update observers for its observable objects. This interface should be implemented by add-ins that wish to observe Ice objects in order to collect statistics. An instance of this interface can be provided to the Ice run-time through the Ice communicator initialization data.
Used By

- Ice::Communicator::getObserver

Operation Index

getConnectionEstablishmentObserver — This method should return an observer for the given endpoint information and connector.
getEndpointLookupObserver — This method should return an observer for the given endpoint information.
getConnectionObserver — This method should return a connection observer for the given connection.
getThreadObserver — This method should return a thread observer for the given thread.
getInvocationObserver — This method should return an invocation observer for the given invocation.
getDispatchObserver — This method should return a dispatch observer for the given dispatch.
setObserverUpdater — The Ice run-time calls this method when the communicator is initialized.

Operations

Ice::Instrumentation::Observer getConnectionEstablishmentObserver(Ice::Endpoint endpt, string connector)

This method should return an observer for the given endpoint information and connector. The Ice run-time calls this method for each connection establishment attempt.
Parameters

`endpt` — The endpoint.
`connector` — The description of the connector. For IP transports, this is typically the IP address to connect to.
Return Value

The observer to instrument the connection establishment.

Ice::Instrumentation::Observer getEndpointLookupObserver(Ice::Endpoint endpt)

This method should return an observer for the given endpoint information. The Ice run-time calls this method to resolve an endpoint and obtain the list of connectors. For IP endpoints, this typically involves doing a DNS lookup to obtain the IP addresses associated with the DNS name.
Parameters

`endpt` — The endpoint.
Return Value

The observer to instrument the endpoint lookup.

Ice::Instrumentation::ConnectionObserver getConnectionObserver(Ice::ConnectionInfo c, Ice::Endpoint e, Ice::Instrumentation::ConnectionState s, Ice::Instrumentation::ConnectionObserver o)

This method should return a connection observer for the given connection. The Ice run-time calls this method for each new connection and for all the Ice communicator connections when Ice::Instrumentation::ObserverUpdater::updateConnectionObservers is called.
Parameters

c — The connection information.
e — The connection endpoint.
s — The state of the connection.
o — The old connection observer if one is already set or a null reference otherwise.
Return Value

The connection observer to instrument the connection.

Copyright 2016, ZeroC, Inc.

Ice::Instrumentation::ThreadObserver getThreadObserver(string parent, string id, Ice::Instrumentation::ThreadState s, Ice::Instrumentation::ThreadObserver o)

This method should return a thread observer for the given thread. The Ice run-time calls this method for each new thread and for all the Ice communicator threads when Ice::Instrumentation::ObserverUpdater::updateThreadObservers is called.
Parameters

`parent` — The parent of the thread.
`id` — The ID of the thread to observe.
`s` — The state of the thread.
`o` — The old thread observer if one is already set or a null reference otherwise.
Return Value

The thread observer to instrument the thread.

Ice::Instrumentation::InvocationObserver getInvocationObserver(Object prx, string operation, Ice::Context ctx)*

This method should return an invocation observer for the given invocation. The Ice run-time calls this method for each new invocation on a proxy.
Parameters

`prx` — The proxy used for the invocation.
`operation` — The name of the invocation.
`ctx` — The context specified by the user.
Return Value

The invocation observer to instrument the invocation.

Ice::Instrumentation::DispatchObserver getDispatchObserver(Ice::Current c, int size)

This method should return a dispatch observer for the given dispatch. The Ice run-time calls this method each time it receives an incoming invocation to be dispatched for an Ice object.
Parameters

`c` — The current object as provided to the Ice servant dispatching the invocation.
`size` — The size of the dispatch.
Return Value

The dispatch observer to instrument the dispatch.

void setObserverUpdater(Ice::Instrumentation::ObserverUpdater updater)

The Ice run-time calls this method when the communicator is initialized. The add-in implementing this interface can use this object to get the Ice run-time to re-obtain observers for observed objects.
Parameters

`updater` — The observer updater object.

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation-ConnectionObserver

Ice::Instrumentation::ConnectionObserver

Overview

local interface ConnectionObserver extends Ice::Instrumentation::Observer

The connection observer interface to instrument Ice connections.
Used By

- Ice::Instrumentation::CommunicatorObserver::getConnectionObserver

Operation Index

sentBytes — Notification of sent bytes over the connection.
receivedBytes — Notification of received bytes over the connection.

Operations

void sentBytes(int num)

Notification of sent bytes over the connection.
Parameters

`num` — The number of bytes sent.

void receivedBytes(int num)

Notification of received bytes over the connection.
Parameters

`num` — The number of bytes received.

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation-ConnectionState

Ice::Instrumentation::ConnectionState

Overview

local enum ConnectionState

The state of an Ice connection.
Used By

- CommunicatorObserver::getConnectionObserver

Enumerator Index

ConnectionStateValidating — The connection is being validated.
ConnectionStateHolding — The connection is holding the reception of new messages.
ConnectionStateActive — The connection is active and can send and receive messages.
ConnectionStateClosing — The connection is being gracefully shutdown and waits for the peer to close its end of the connection.
ConnectionStateClosed — The connection is closed and waits for potential dispatch to be finished before being destroyed and detached from the observer.

Enumerators

ConnectionStateValidating

The connection is being validated.

ConnectionStateHolding

The connection is holding the reception of new messages.

ConnectionStateActive

The connection is active and can send and receive messages.

ConnectionStateClosing

The connection is being gracefully shutdown and waits for the peer to close its end of the connection.

ConnectionStateClosed

The connection is closed and waits for potential dispatch to be finished before being destroyed and detached from the observer.

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation-DispatchObserver

Ice::Instrumentation::DispatchObserver

Overview

local interface DispatchObserver extends Ice::Instrumentation::Observer

The dispatch observer to instrument servant dispatch.
Used By

- Ice::Instrumentation::CommunicatorObserver::getDispatchObserver

Operation Index

userException — Notification of a user exception.
reply — Reply notification.

Operations

void userException()

Notification of a user exception.

void reply(int size)

Reply notification.
Parameters

`size` — The size of the reply.

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation-InvocationObserver

Ice::Instrumentation::InvocationObserver

Overview

local interface InvocationObserver extends Ice::Instrumentation::Observer

The invocation observer to instrument invocations on proxies. A proxy invocation can either result in a collocated or remote invocation. If it results in a remote invocation, a sub-observer is requested for the remote invocation.
Used By

- Ice::Instrumentation::CommunicatorObserver::getInvocationObserver

Operation Index

retried — Notification of the invocation being retried.
userException — Notification of a user exception.
getRemoteObserver — Get a remote observer for this invocation.
getCollocatedObserver — Get a collocated observer for this invocation.

Operations

void retried()

Notification of the invocation being retried.

void userException()

Notification of a user exception.

Ice::Instrumentation::RemoteObserver getRemoteObserver(Ice::ConnectionInfo con, Ice::Endpoint endpt, int requestId, int size)

Get a remote observer for this invocation.
Parameters

`con` — The connection information.
`endpt` — The connection endpoint.
`requestId` — The ID of the invocation.
`size` — The size of the invocation.
Return Value

The observer to instrument the remote invocation.

Ice::Instrumentation::CollocatedObserver getCollocatedObserver(Ice::ObjectAdapter adapter, int requestId, int size)

Get a collocated observer for this invocation.
Parameters

`adapter` — The object adapter hosting the collocated Ice object.
`requestId` — The ID of the invocation.
`size` — The size of the invocation.
Return Value

The observer to instrument the collocated invocation.

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation-Observer

Ice::Instrumentation::Observer

Overview

local interface Observer

The object observer interface used by instrumented objects to notify the observer of their existence.
Derived Classes and Interfaces

- Ice::Instrumentation::ChildInvocationObserver
- Ice::Instrumentation::ConnectionObserver
- Ice::Instrumentation::DispatchObserver
- Ice::Instrumentation::InvocationObserver
- Ice::Instrumentation::ThreadObserver

Used By

- Ice::Instrumentation::CommunicatorObserver::getConnectionEstablishmentObserver
- Ice::Instrumentation::CommunicatorObserver::getEndpointLookupObserver

Operation Index

attach — This method is called when the instrumented object is created or when the observer is attached to an existing object.
detach — This method is called when the instrumented object is destroyed and as a result the observer detached from the object.
failed — Notification of a failure.

Operations

void attach()

This method is called when the instrumented object is created or when the observer is attached to an existing object.

void detach()

This method is called when the instrumented object is destroyed and as a result the observer detached from the object.

void failed(string exceptionName)

Notification of a failure.
Parameters

`exceptionName` — The name of the exception.

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation-ObserverUpdater

Ice::Instrumentation::ObserverUpdater

Overview

local interface ObserverUpdater

The observer updater interface. This interface is implemented by the Ice run-time and an instance of this interface is provided by the Ice communicator on initialization to the Ice::Instrumentation::CommunicatorObserver object set with the communicator initialization data. The Ice communicator calls Ice::Instrumentation::CommunicatorObserver::setObserverUpdater to provide the observer updater. This interface can be used by add-ins implementing the Ice::Instrumentation::CommunicatorObserver interface to update the observers of connections and threads.
Used By

- Ice::Instrumentation::CommunicatorObserver::setObserverUpdater

Operation Index

updateConnectionObservers — Update connection observers associated with each of the Ice connection from the communicator and its object adapters.
updateThreadObservers — Update thread observers associated with each of the Ice thread from the communicator and its object adapters.

Operations

void updateConnectionObservers()

Update connection observers associated with each of the Ice connection from the communicator and its object adapters. When called, this method goes through all the connections and for each connection Ice::Instrumentation::CommunicatorObserver::getConnectionObserver is called. The implementation of getConnectionObserver has the possibility to return an updated observer if necessary.

void updateThreadObservers()

Update thread observers associated with each of the Ice thread from the communicator and its object adapters. When called, this method goes through all the threads and for each thread Ice::Instrumentation::CommunicatorObserver::getThreadObserver is called. The implementation of getThreadObserver has the possibility to return an updated observer if necessary.

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation-RemoteObserver

Ice::Instrumentation::RemoteObserver

Overview

local interface RemoteObserver extends Ice::Instrumentation::ChildInvocationObserver

The remote observer to instrument invocations that are sent over the wire.
Used By

- Ice::Instrumentation::InvocationObserver::getRemoteObserver

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation-ThreadObserver

Ice::Instrumentation::ThreadObserver

Overview

local interface ThreadObserver extends Ice::Instrumentation::Observer

The thread observer interface to instrument Ice threads. This can be threads from the Ice thread pool or utility threads used by the Ice core.
Used By

- Ice::Instrumentation::CommunicatorObserver::getThreadObserver

Operation Index

stateChanged — Notification of thread state change.

Operations

void stateChanged(Ice::Instrumentation::ThreadState oldState, Ice::Instrumentation::ThreadState newState)

Notification of thread state change.
Parameters

`oldState` — The previous thread state.
`newState` — The new thread state.

Copyright 2016, ZeroC, Inc.

Ice-Instrumentation-ThreadState

Ice::Instrumentation::ThreadState

Overview

local enum ThreadState

The thread state enumeration keeps track of the different possible states of Ice threads.
Used By

- CommunicatorObserver::getThreadObserver
- ThreadObserver::stateChanged

Enumerator Index

ThreadStateIdle — The thread is idle.
ThreadStateInUseForIO — The thread is in use performing reads or writes for Ice connections.
ThreadStateInUseForUser — The thread is calling user code (servant implementation, AMI callbacks).
ThreadStateInUseForOther — The thread is performing other internal activities (DNS lookups, timer callbacks, etc).

Enumerators

ThreadStateIdle

The thread is idle.

ThreadStateInUseForIO

The thread is in use performing reads or writes for Ice connections. This state is only for threads from an Ice thread pool.

ThreadStateInUseForUser

The thread is calling user code (servant implementation, AMI callbacks). This state is only for threads from an Ice thread pool.

ThreadStateInUseForOther

The thread is performing other internal activities (DNS lookups, timer callbacks, etc).

Copyright 2016, ZeroC, Inc.

Ice-InvalidReplicaGroupIdException

Ice::InvalidReplicaGroupIdException

Overview

exception InvalidReplicaGroupIdException

This exception is raised if the replica group provided by the server is invalid.

Used By

- Ice::LocatorRegistry::setReplicatedAdapterDirectProxy

Copyright 2016, ZeroC, Inc.

Ice-InvocationCanceledException

Ice::InvocationCanceledException

Overview

local exception InvocationCanceledException

This exception indicates that an asynchronous invocation failed because it was canceled explicitly by the user using the `Ice::AsyncResult::cancel` method.

Copyright 2016, ZeroC, Inc.

Ice-InvocationTimeoutException

Ice::InvocationTimeoutException

Overview

local exception InvocationTimeoutException extends Ice::TimeoutException

This exception indicates that an invocation failed because it timed out.

Copyright 2016, ZeroC, Inc.

Ice-IPConnectionInfo

Ice::IPConnectionInfo

Overview

local class IPConnectionInfo extends Ice::ConnectionInfo

Provides access to the connection details of an IP connection

Derived Classes and Interfaces

- Ice::TCPConnectionInfo
- Ice::UDPConnectionInfo
- IceSSL::ConnectionInfo

Data Member Index

localAddress — The local address.
localPort — The local port.
remoteAddress — The remote address.
remotePort — The remote port.

Data Members

string localAddress;

The local address.

int localPort;

The local port.

string remoteAddress;

The remote address.

int remotePort;

The remote port.

Copyright 2016, ZeroC, Inc.

Ice-IPEndpointInfo

Ice::IPEndpointInfo

Overview

local class IPEndpointInfo extends **Ice::EndpointInfo**

Provides access to the address details of a IP endpoint.

Derived Classes and Interfaces

- Ice::TCPEndpointInfo
- Ice::UDPEndpointInfo
- IceSSL::EndpointInfo

See Also

- Ice::Endpoint

Data Member Index

host — The host or address configured with the endpoint.
port — The port number.
sourceAddress — The source IP address.

Data Members

string host;

The host or address configured with the endpoint.

int port;

The port number.

string sourceAddress;

The source IP address.

Copyright 2016, ZeroC, Inc.

Ice-Locator

Ice::Locator

Overview

interface Locator

The Ice locator interface. This interface is used by clients to lookup adapters and objects. It is also used by servers to get the locator registry proxy.

The Ice::Locator interface is intended to be used by Ice internals and by locator implementations. Regular user code should not attempt to use any functionality of this interface directly.

Derived Classes and Interfaces

- IceGrid::Locator

Operation Index

findObjectById — Find an object by identity and return a proxy that contains the adapter ID or endpoints which can be used to access the object.
findAdapterById — Find an adapter by id and return a proxy that contains its endpoints.
getRegistry — Get the locator registry.

Operations

["amd", "nonmutating"] Object* findObjectById(**Ice::Identity** id) throws **Ice::ObjectNotFoundException**

Find an object by identity and return a proxy that contains the adapter ID or endpoints which can be used to access the object.

Parameters

`id` — The identity.

Return Value

The proxy, or null if the object is not active.

Exceptions

Ice::ObjectNotFoundException — Raised if the object cannot be found.

["amd", "nonmutating"] Object* findAdapterById(string id) throws **Ice::AdapterNotFoundException**

Find an adapter by id and return a proxy that contains its endpoints.

Parameters

`id` — The adapter id.

Return Value

The adapter proxy, or null if the adapter is not active.

Exceptions

Ice::AdapterNotFoundException — Raised if the adapter cannot be found.

["nonmutating"] **Ice::LocatorRegistry*** getRegistry()

Get the locator registry.

Copyright 2016, ZeroC, Inc.

Return Value

The locator registry.

Copyright 2016, ZeroC, Inc.

Ice-LocatorFinder

Ice::LocatorFinder

Overview

interface LocatorFinder

This inferface should be implemented by services implementing the Ice::Locator interface. It should be advertised through an Ice object with the identity `Ice/LocatorFinder'. This allows clients to retrieve the locator proxy with just the endpoint information of the service.

Operation Index

getLocator — Get the locator proxy implemented by the process hosting this finder object.

Operations

Ice::Locator* getLocator()

Get the locator proxy implemented by the process hosting this finder object. The proxy might point to several replicas.

Return Value

The locator proxy.

Copyright 2016, ZeroC, Inc.

Ice-LocatorRegistry

Ice::LocatorRegistry

Overview

interface LocatorRegistry

The Ice locator registry interface. This interface is used by servers to register adapter endpoints with the locator.

The Ice::LocatorRegistry interface is intended to be used by Ice internals and by locator implementations. Regular user code should not attempt to use any functionality of this interface directly.

Operation Index

setAdapterDirectProxy — Set the adapter endpoints with the locator registry.
setReplicatedAdapterDirectProxy — Set the adapter endpoints with the locator registry.
setServerProcessProxy — Set the process proxy for a server.

Operations

["amd"] void setAdapterDirectProxy(string id, Object* proxy) throws Ice::AdapterNotFoundException, Ice::AdapterAlreadyActiveException

Set the adapter endpoints with the locator registry.

Parameters

`id` — The adapter id.
`proxy` — The adapter proxy (a dummy direct proxy created by the adapter). The direct proxy contains the adapter endpoints.

Exceptions

Ice::AdapterNotFoundException — Raised if the adapter cannot be found, or if the locator only allows registered adapters to set their active proxy and the adapter is not registered with the locator.
Ice::AdapterAlreadyActiveException — Raised if an adapter with the same id is already active.

["amd"] void setReplicatedAdapterDirectProxy(string adapterId, string replicaGroupId, Object* p) throws Ice::AdapterNotFoundException, Ice::AdapterAlreadyActiveException, Ice::InvalidReplicaGroupIdException

Set the adapter endpoints with the locator registry.

Parameters

`adapterId` — The adapter id.
`replicaGroupId` — The replica group id.
p — The adapter proxy (a dummy direct proxy created by the adapter). The direct proxy contains the adapter endpoints.

Exceptions

Ice::AdapterNotFoundException — Raised if the adapter cannot be found, or if the locator only allows registered adapters to set their active proxy and the adapter is not registered with the locator.
Ice::AdapterAlreadyActiveException — Raised if an adapter with the same id is already active.
Ice::InvalidReplicaGroupIdException — Raised if the given replica group doesn't match the one registered with the locator registry for this object adapter.

["amd"] void setServerProcessProxy(string id, Ice::Process* proxy) throws Ice::ServerNotFoundException

Set the process proxy for a server.

Parameters

Copyright 2016, ZeroC, Inc.

`id` — The server id.
`proxy` — The process proxy.

Exceptions

Ice::ServerNotFoundException — Raised if the server cannot be found.

Copyright 2016, ZeroC, Inc.

Ice-Logger

Ice::Logger

Overview

local interface Logger

The Ice message logger. Applications can provide their own logger by implementing this interface and installing it in a communicator.

Used By

- Ice::Communicator::getLogger
- cloneWithPrefix

Operation Index

print — Print a message.
trace — Log a trace message.
warning — Log a warning message.
error — Log an error message.
getPrefix — Returns this logger's prefix.
cloneWithPrefix — Returns a clone of the logger with a new prefix.

Operations

void print(string message)

Print a message. The message is printed literally, without any decorations such as executable name or time stamp.

void trace(string category, string message)

Log a trace message.

Parameters

`category` — The trace category.
`message` — The trace message to log.

void warning(string message)

Log a warning message.

Parameters

`message` — The warning message to log.

See Also

- error

void error(string message)

Log an error message.

Parameters

`message` — The error message to log.

See Also

- warning

Copyright 2016, ZeroC, Inc.

string getPrefix()

Returns this logger's prefix.

Return Value

The prefix.

Ice::Logger cloneWithPrefix(string prefix)

Returns a clone of the logger with a new prefix.

Parameters

`prefix` — The new prefix for the logger.

Copyright 2016, ZeroC, Inc.

Ice-LoggerAdmin

Ice::LoggerAdmin

Overview

interface LoggerAdmin

The interface of the admin object that allows an Ice application the attach its Ice::RemoteLogger to the Ice::Logger of this admin object's Ice communicator.

Operation Index

attachRemoteLogger — attachRemoteLogger is called to attach a Ice::RemoteLogger object to the local Ice::Logger.
detachRemoteLogger — detachRemoteLogger is called to detach a Ice::RemoteLogger object from the local Ice::Logger.
getLog — getLog retrieves log messages recently logged.

Operations

void attachRemoteLogger(**Ice::RemoteLogger*** prx, **Ice::LogMessageTypeSeq** messageTypes, **Ice::StringSeq** traceCategories, int messageMax) throws **Ice::RemoteLoggerAlreadyAttachedException**

attachRemoteLogger is called to attach a Ice::RemoteLogger object to the local Ice::Logger. attachRemoteLogger calls init on the provided Ice::RemoteLogger proxy.

Parameters

`prx` — A proxy to the remote logger.
`messageTypes` — The list of message types that the remote logger wishes to receive. An empty list means no filtering (send all message types).
`traceCategories` — The categories of traces that the remote logger wishes to receive. This parameter is ignored if messageTypes is not empty and does not include trace. An empty list means no filtering (send all trace categories).
`messageMax` — The maximum number of log messages (of all types) to be provided to Ice::RemoteLogger::init. A negative value requests all messages available.

Exceptions

Ice::RemoteLoggerAlreadyAttachedException — Raised if this remote logger is already attached to this admin object.

bool detachRemoteLogger(**Ice::RemoteLogger*** prx)

detachRemoteLogger is called to detach a Ice::RemoteLogger object from the local Ice::Logger.

Parameters

`prx` — A proxy to the remote logger.

Return Value

True if the provided remote logger proxy was detached, and false otherwise.

Ice::LogMessageSeq getLog(**Ice::LogMessageTypeSeq** messageTypes, **Ice::StringSeq** traceCategories, int messageMax, out string prefix)

getLog retrieves log messages recently logged.

Parameters

`messageTypes` — The list of message types that the caller wishes to receive. An empty list means no filtering (send all message types).
`traceCategories` — The categories of traces that caller wish to receive. This parameter is ignored if messageTypes is not empty and

Copyright 2016, ZeroC, Inc.

does not include trace. An empty list means no filtering (send all trace categories).
`messageMax` — The maximum number of log messages (of all types) to be returned. A negative value requests all messages available.
`prefix` — The prefix of the associated local Ice::Logger.

Return Value

The Log messages.

Copyright 2016, ZeroC, Inc.

Ice-LogMessage

Ice::LogMessage

Overview

struct LogMessage

A complete log message.

Used By

- Ice::LogMessageSeq
- Ice::RemoteLogger::log

Data Member Index

type — The type of message sent to the Ice::Logger.
timestamp — The date and time when the Ice::Logger received this message, expressed as the number of microseconds since the Unix Epoch (00:00:00 UTC on 1 January 1970)
traceCategory — For a message of type trace, the trace category of this log message; otherwise, the empty string.
message — The log message itself.

Data Members

Ice::LogMessageType type;

The type of message sent to the Ice::Logger.

long timestamp;

The date and time when the Ice::Logger received this message, expressed as the number of microseconds since the Unix Epoch (00:00:00 UTC on 1 January 1970)

string traceCategory;

For a message of type trace, the trace category of this log message; otherwise, the empty string.

string message;

The log message itself.

Copyright 2016, ZeroC, Inc.

Ice-LogMessageType

Ice::LogMessageType

Overview

enum LogMessageType

An enumeration representing the different types of log messages.

Used By

- LogMessage::type
- LogMessageTypeSeq

Enumerator Index

PrintMessage — The Logger received a print message.
TraceMessage — The Logger received a trace message.
WarningMessage — The Logger received a warning message.
ErrorMessage — The Logger received an error message.

Enumerators

PrintMessage

The Logger received a print message.

TraceMessage

The Logger received a trace message.

WarningMessage

The Logger received a warning message.

ErrorMessage

The Logger received an error message.

Copyright 2016, ZeroC, Inc.

Ice-MarshalException

Ice::MarshalException

Overview

local exception MarshalException extends Ice::ProtocolException

This exception is raised for errors during marshaling or unmarshaling data.

Derived Exceptions

- Ice::EncapsulationException
- Ice::MemoryLimitException
- Ice::NoObjectFactoryException
- Ice::ProxyUnmarshalException
- Ice::StringConversionException
- Ice::UnexpectedObjectException
- Ice::UnmarshalOutOfBoundsException

Copyright 2016, ZeroC, Inc.

Ice-MemoryLimitException

Ice::MemoryLimitException

Overview

local exception MemoryLimitException extends Ice::MarshalException

This exception is raised when Ice receives a request or reply message whose size exceeds the limit specified by the `Ice.MessageSizeMax` property.

Copyright 2016, ZeroC, Inc.

Ice-NoEndpointException

Ice::NoEndpointException

Overview

local exception NoEndpointException

This exception is raised if no suitable endpoint is available.

Data Member Index

proxy — The stringified proxy for which no suitable endpoint is available.

Data Members

string proxy;

The stringified proxy for which no suitable endpoint is available.

Copyright 2016, ZeroC, Inc.

Ice-NoObjectFactoryException

Ice::NoObjectFactoryException

Overview

local exception NoObjectFactoryException extends Ice::MarshalException

This exception is raised if no suitable object factory was found during unmarshaling of a Slice class instance.

See Also

- Ice::ObjectFactory
- Ice::Communicator::addObjectFactory
- Ice::Communicator::findObjectFactory

Data Member Index

type — The Slice type ID of the class instance for which no no factory could be found.

Data Members

string type;

The Slice type ID of the class instance for which no no factory could be found.

Copyright 2016, ZeroC, Inc.

Ice-NotRegisteredException

Ice::NotRegisteredException

Overview

local exception NotRegisteredException

An attempt was made to find or deregister something that is not registered with the Ice run time or Ice locator. This exception is raised if an attempt is made to remove a servant, servant locator, facet, object factory, plug-in, object adapter, object, or user exception factory that is not currently registered. It's also raised if the Ice locator can't find an object or object adapter when resolving an indirect proxy or when an object adapter is activated.

Data Member Index

kindOfObject — The kind of object that could not be removed: "servant", "servant locator", "object factory", "plug-in", "object adapter", "object", or "user exception factory".
id — The ID (or name) of the object that could not be removed.

Data Members

string kindOfObject;

The kind of object that could not be removed: "servant", "servant locator", "object factory", "plug-in", "object adapter", "object", or "user exception factory".

string id;

The ID (or name) of the object that could not be removed.

Copyright 2016, ZeroC, Inc.

Ice-ObjectAdapter

Ice::ObjectAdapter

Overview

local interface ObjectAdapter

The object adapter provides an up-call interface from the Ice run time to the implementation of Ice objects. The object adapter is responsible for receiving requests from endpoints, and for mapping between servants, identities, and proxies.

Used By

- Freeze::ServantInitializer::initialize
- Ice::Communicator::createAdmin
- Ice::Communicator::createObjectAdapterWithEndpoints
- Ice::Communicator::createObjectAdapterWithRouter
- Ice::Communicator::createObjectAdapter
- Ice::Connection::getAdapter
- Ice::Connection::setAdapter
- Ice::Current::adapter
- Ice::Instrumentation::InvocationObserver::getCollocatedObserver

See Also

- Ice::Communicator
- Ice::ServantLocator

Operation Index

getName — Get the name of this object adapter.
getCommunicator — Get the communicator this object adapter belongs to.
activate — Activate all endpoints that belong to this object adapter.
hold — Temporarily hold receiving and dispatching requests.
waitForHold — Wait until the object adapter holds requests.
deactivate — Deactivate all endpoints that belong to this object adapter.
waitForDeactivate — Wait until the object adapter has deactivated.
isDeactivated — Check whether object adapter has been deactivated.
destroy — Destroys the object adapter and cleans up all resources held by the object adapter.
add — Add a servant to this object adapter's Active Servant Map.
addFacet — Like add, but with a facet.
addWithUUID — Add a servant to this object adapter's Active Servant Map, using an automatically generated UUID as its identity.
addFacetWithUUID — Like addWithUUID, but with a facet.
addDefaultServant — Add a default servant to handle requests for a specific category.
remove — Remove a servant (that is, the default facet) from the object adapter's Active Servant Map.
removeFacet — Like remove, but with a facet.
removeAllFacets — Remove all facets with the given identity from the Active Servant Map.
removeDefaultServant — Remove the default servant for a specific category.
find — Look up a servant in this object adapter's Active Servant Map by the identity of the Ice object it implements.
findFacet — Like find, but with a facet.
findAllFacets — Find all facets with the given identity in the Active Servant Map.
findByProxy — Look up a servant in this object adapter's Active Servant Map, given a proxy.
addServantLocator — Add a Servant Locator to this object adapter.
removeServantLocator — Remove a Servant Locator from this object adapter.
findServantLocator — Find a Servant Locator installed with this object adapter.
findDefaultServant — Find the default servant for a specific category.
createProxy — Create a proxy for the object with the given identity.
createDirectProxy — Create a direct proxy for the object with the given identity.
createIndirectProxy — Create an indirect proxy for the object with the given identity.
setLocator — Set an Ice locator for this object adapter.
getLocator — Get the Ice locator used by this object adapter.
refreshPublishedEndpoints — Refresh the set of published endpoints.
getEndpoints — Get the set of endpoints configured with this object adapter.
getPublishedEndpoints — Get the set of endpoints that proxies created by this object adapter will contain.

Copyright 2016, ZeroC, Inc.

Operations

string getName()

Get the name of this object adapter.

Return Value

This object adapter's name.

Ice::Communicator getCommunicator()

Get the communicator this object adapter belongs to.

Return Value

This object adapter's communicator.

See Also

- Ice::Communicator

void activate()

Activate all endpoints that belong to this object adapter. After activation, the object adapter can dispatch requests received through its endpoints.

See Also

- hold
- deactivate

void hold()

Temporarily hold receiving and dispatching requests. The object adapter can be reactivated with the activate operation.

Holding is not immediate, i.e., after hold returns, the object adapter might still be active for some time. You can use waitForHold to wait until holding is complete.

See Also

- activate
- deactivate
- waitForHold

void waitForHold()

Wait until the object adapter holds requests. Calling hold initiates holding of requests, and waitForHold only returns when holding of requests has been completed.

See Also

- hold
- waitForDeactivate
- Ice::Communicator::waitForShutdown

void deactivate()

Deactivate all endpoints that belong to this object adapter. After deactivation, the object adapter stops receiving requests through its endpoints. Object adapters that have been deactivated must not be reactivated again, and cannot be used otherwise. Attempts to use a deactivated object adapter raise Ice::ObjectAdapterDeactivatedException; however, attempts to deactivate an already deactivated object adapter are ignored and do nothing. Once deactivated, it is possible to destroy the adapter to clean up resources and then create and activate a new adapter with the same name.

After deactivate returns, no new requests are processed by the object adapter. However, requests that have been started before deactivate was called might still be active. You can use waitForDeactivate to wait for the completion of all requests for this object adapter.

Copyright 2016, ZeroC, Inc.

See Also

- activate
- hold
- waitForDeactivate
- Ice::Communicator::shutdown

void waitForDeactivate()

Wait until the object adapter has deactivated. Calling deactivate initiates object adapter deactivation, and waitForDeactivate only returns when deactivation has been completed.

See Also

- deactivate
- waitForHold
- Ice::Communicator::waitForShutdown

bool isDeactivated()

Check whether object adapter has been deactivated.

Return Value

Whether adapter has been deactivated.

See Also

- Ice::Communicator::shutdown

void destroy()

Destroys the object adapter and cleans up all resources held by the object adapter. If the object adapter has not yet been deactivated, destroy implicitly initiates the deactivation and waits for it to finish. Subsequent calls to destroy are ignored. Once destroy has returned, it is possible to create another object adapter with the same name.

See Also

- deactivate
- waitForDeactivate
- Ice::Communicator::destroy

Object* add(Object servant, **Ice::Identity** id)

Add a servant to this object adapter's Active Servant Map. Note that one servant can implement several Ice objects by registering the servant with multiple identities. Adding a servant with an identity that is in the map already throws Ice::AlreadyRegisteredException.

Parameters

`servant` — The servant to add.
`id` — The identity of the Ice object that is implemented by the servant.

Return Value

A proxy that matches the given identity and this object adapter.

See Also

- Ice::Identity
- addFacet
- addWithUUID
- remove
- find

Object* addFacet(Object servant, **Ice::Identity** id, string facet)

Copyright 2016, ZeroC, Inc.

Like add, but with a facet. Calling `add(servant, id)` is equivalent to calling addFacet with an empty facet.

Parameters

`servant` — The servant to add.
`id` — The identity of the Ice object that is implemented by the servant.
`facet` — The facet. An empty facet means the default facet.

Return Value

A proxy that matches the given identity, facet, and this object adapter.

See Also

- Ice::Identity
- add
- addFacetWithUUID
- removeFacet
- findFacet

Object* addWithUUID(Object servant)

Add a servant to this object adapter's Active Servant Map, using an automatically generated UUID as its identity. Note that the generated UUID identity can be accessed using the proxy's `ice_getIdentity` operation.

Parameters

`servant` — The servant to add.

Return Value

A proxy that matches the generated UUID identity and this object adapter.

See Also

- Ice::Identity
- add
- addFacetWithUUID
- remove
- find

Object* addFacetWithUUID(Object servant, string facet)

Like addWithUUID, but with a facet. Calling `addWithUUID(servant)` is equivalent to calling addFacetWithUUID with an empty facet.

Parameters

`servant` — The servant to add.
`facet` — The facet. An empty facet means the default facet.

Return Value

A proxy that matches the generated UUID identity, facet, and this object adapter.

See Also

- Ice::Identity
- addFacet
- addWithUUID
- removeFacet
- findFacet

void addDefaultServant(Object servant, string category)

Add a default servant to handle requests for a specific category. Adding a default servant for a category for which a default servant is already

Copyright 2016, ZeroC, Inc.

registered throws Ice::AlreadyRegisteredException. To dispatch operation calls on servants, the object adapter tries to find a servant for a given Ice object identity and facet in the following order:

1. The object adapter tries to find a servant for the identity and facet in the Active Servant Map.
2. If no servant has been found in the Active Servant Map, the object adapter tries to find a default servant for the category component of the identity.
3. If no servant has been found by any of the preceding steps, the object adapter tries to find a default servant for an empty category, regardless of the category contained in the identity.
4. If no servant has been found by any of the preceding steps, the object adapter gives up and the caller receives Ice::ObjectNotExistException or Ice::FacetNotExistException.

Parameters

`servant` — The default servant.
`category` — The category for which the default servant is registered. An empty category means it will handle all categories.

See Also

- removeDefaultServant
- findDefaultServant

Object remove(**Ice::Identity** id)

Remove a servant (that is, the default facet) from the object adapter's Active Servant Map.

Parameters

`id` — The identity of the Ice object that is implemented by the servant. If the servant implements multiple Ice objects, remove has to be called for all those Ice objects. Removing an identity that is not in the map throws Ice::NotRegisteredException.

Return Value

The removed servant.

See Also

- Ice::Identity
- add
- addWithUUID

Object removeFacet(**Ice::Identity** id, string facet)

Like remove, but with a facet. Calling `remove(id)` is equivalent to calling removeFacet with an empty facet.

Parameters

`id` — The identity of the Ice object that is implemented by the servant.
`facet` — The facet. An empty facet means the default facet.

Return Value

The removed servant.

See Also

- Ice::Identity
- addFacet
- addFacetWithUUID

Ice::FacetMap removeAllFacets(**Ice::Identity** id)

Remove all facets with the given identity from the Active Servant Map. The operation completely removes the Ice object, including its default facet. Removing an identity that is not in the map throws Ice::NotRegisteredException.

Parameters

Copyright 2016, ZeroC, Inc.

`id` — The identity of the Ice object to be removed.

Return Value

A collection containing all the facet names and servants of the removed Ice object.

See Also

- remove
- removeFacet

Object removeDefaultServant(string category)

Remove the default servant for a specific category. Attempting to remove a default servant for a category that is not registered throws Ice::NotRegisteredException.

Parameters

`category` — The category of the default servant to remove.

Return Value

The default servant.

See Also

- addDefaultServant
- findDefaultServant

Object find(**Ice::Identity** id)

Look up a servant in this object adapter's Active Servant Map by the identity of the Ice object it implements.

This operation only tries to look up a servant in the Active Servant Map. It does not attempt to find a servant by using any installed Ice::ServantLocator.

Parameters

`id` — The identity of the Ice object for which the servant should be returned.

Return Value

The servant that implements the Ice object with the given identity, or null if no such servant has been found.

See Also

- Ice::Identity
- findFacet
- findByProxy

Object findFacet(**Ice::Identity** id, string facet)

Like find, but with a facet. Calling `find(id)` is equivalent to calling findFacet with an empty facet.

Parameters

`id` — The identity of the Ice object for which the servant should be returned.
`facet` — The facet. An empty facet means the default facet.

Return Value

The servant that implements the Ice object with the given identity and facet, or null if no such servant has been found.

See Also

- Ice::Identity

Copyright 2016, ZeroC, Inc.

- find
- findByProxy

Ice::FacetMap findAllFacets(Ice::Identity id)

Find all facets with the given identity in the Active Servant Map.

Parameters

`id` — The identity of the Ice object for which the facets should be returned.

Return Value

A collection containing all the facet names and servants that have been found, or an empty map if there is no facet for the given identity.

See Also

- find
- findFacet

Object findByProxy(Object* proxy)

Look up a servant in this object adapter's Active Servant Map, given a proxy.

This operation only tries to lookup a servant in the Active Servant Map. It does not attempt to find a servant by using any installed Ice::ServantLocator.

Parameters

`proxy` — The proxy for which the servant should be returned.

Return Value

The servant that matches the proxy, or null if no such servant has been found.

See Also

- find
- findFacet

void addServantLocator(Ice::ServantLocator locator, string category)

Add a Servant Locator to this object adapter. Adding a servant locator for a category for which a servant locator is already registered throws Ice::AlreadyRegisteredException. To dispatch operation calls on servants, the object adapter tries to find a servant for a given Ice object identity and facet in the following order:

1. The object adapter tries to find a servant for the identity and facet in the Active Servant Map.
2. If no servant has been found in the Active Servant Map, the object adapter tries to find a servant locator for the category component of the identity. If a locator is found, the object adapter tries to find a servant using this locator.
3. If no servant has been found by any of the preceding steps, the object adapter tries to find a locator for an empty category, regardless of the category contained in the identity. If a locator is found, the object adapter tries to find a servant using this locator.
4. If no servant has been found by any of the preceding steps, the object adapter gives up and the caller receives Ice::ObjectNotExistException or Ice::FacetNotExistException.

Only one locator for the empty category can be installed.

Parameters

`locator` — The locator to add.
`category` — The category for which the Servant Locator can locate servants, or an empty string if the Servant Locator does not belong to any specific category.

See Also

- Ice::Identity
- removeServantLocator
- findServantLocator

Copyright 2016, ZeroC, Inc.

- Ice::ServantLocator

Ice::ServantLocator removeServantLocator(string category)

Remove a Servant Locator from this object adapter.

Parameters

`category` — The category for which the Servant Locator can locate servants, or an empty string if the Servant Locator does not belong to any specific category.

Return Value

The Servant Locator, or throws Ice::NotRegisteredException if no Servant Locator was found for the given category.

See Also

- Ice::Identity
- addServantLocator
- findServantLocator
- Ice::ServantLocator

Ice::ServantLocator findServantLocator(string category)

Find a Servant Locator installed with this object adapter.

Parameters

`category` — The category for which the Servant Locator can locate servants, or an empty string if the Servant Locator does not belong to any specific category.

Return Value

The Servant Locator, or null if no Servant Locator was found for the given category.

See Also

- Ice::Identity
- addServantLocator
- removeServantLocator
- Ice::ServantLocator

Object findDefaultServant(string category)

Find the default servant for a specific category.

Parameters

`category` — The category of the default servant to find.

Return Value

The default servant or null if no default servant was registered for the category.

See Also

- addDefaultServant
- removeDefaultServant

Object* createProxy(Ice::Identity id)

Create a proxy for the object with the given identity. If this object adapter is configured with an adapter id, the return value is an indirect proxy that refers to the adapter id. If a replica group id is also defined, the return value is an indirect proxy that refers to the replica group id. Otherwise, if no adapter id is defined, the return value is a direct proxy containing this object adapter's published endpoints.

Parameters

Copyright 2016, ZeroC, Inc.

`id` — The object's identity.

Return Value

A proxy for the object with the given identity.

See Also

- Ice::Identity

Object* createDirectProxy(Ice::Identity id)

Create a direct proxy for the object with the given identity. The returned proxy contains this object adapter's published endpoints.

Parameters

`id` — The object's identity.

Return Value

A proxy for the object with the given identity.

See Also

- Ice::Identity

Object* createIndirectProxy(Ice::Identity id)

Create an indirect proxy for the object with the given identity. If this object adapter is configured with an adapter id, the return value refers to the adapter id. Otherwise, the return value contains only the object identity.

Parameters

`id` — The object's identity.

Return Value

A proxy for the object with the given identity.

See Also

- Ice::Identity

void setLocator(Ice::Locator* loc)

Set an Ice locator for this object adapter. By doing so, the object adapter will register itself with the locator registry when it is activated for the first time. Furthermore, the proxies created by this object adapter will contain the adapter identifier instead of its endpoints. The adapter identifier must be configured using the AdapterId property.

Parameters

`loc` — The locator used by this object adapter.

See Also

- createDirectProxy
- Ice::Locator
- Ice::LocatorRegistry

Ice::Locator* getLocator()

Get the Ice locator used by this object adapter.

Return Value

Copyright 2016, ZeroC, Inc.

The locator used by this object adapter, or null if no locator is used by this object adapter.

See Also

- Ice::Locator
- setLocator

void refreshPublishedEndpoints()

Refresh the set of published endpoints. The run time re-reads the PublishedEndpoints property if it is set and re-reads the list of local interfaces if the adapter is configured to listen on all endpoints. This operation is useful to refresh the endpoint information that is published in the proxies that are created by an object adapter if the network interfaces used by a host changes.

Ice::EndpointSeq getEndpoints()

Get the set of endpoints configured with this object adapter.

Return Value

The set of endpoints.

See Also

- Ice::Endpoint

Ice::EndpointSeq getPublishedEndpoints()

Get the set of endpoints that proxies created by this object adapter will contain.

Return Value

The set of published endpoints.

See Also

- refreshPublishedEndpoints
- Ice::Endpoint

Copyright 2016, ZeroC, Inc.

Ice-ObjectAdapterDeactivatedException

Ice::ObjectAdapterDeactivatedException

Overview

local exception ObjectAdapterDeactivatedException

This exception is raised if an attempt is made to use a deactivated Ice::ObjectAdapter.

See Also

- Ice::ObjectAdapter::deactivate
- Ice::Communicator::shutdown

Data Member Index

name — Name of the adapter.

Data Members

string name;

Name of the adapter.

Copyright 2016, ZeroC, Inc.

Ice-ObjectAdapterIdInUseException

Ice::ObjectAdapterIdInUseException

Overview

local exception ObjectAdapterIdInUseException

This exception is raised if an Ice::ObjectAdapter cannot be activated. This happens if the Ice::Locator detects another active Ice::ObjectAdapter with the same adapter id.

Data Member Index

id — Adapter ID.

Data Members

string id;

Adapter ID.

Copyright 2016, ZeroC, Inc.

Ice-ObjectFactory

Ice::ObjectFactory

Overview

local interface ObjectFactory

A factory for objects. Object factories are used in several places, for example, when receiving "objects by value" and when Freeze restores a persistent object. Object factories must be implemented by the application writer, and registered with the communicator.

Used By

- Ice::Communicator::addObjectFactory
- Ice::Communicator::findObjectFactory

Operation Index

create — Create a new object for a given object type.
destroy — Called when the factory is removed from the communicator, or if the communicator is destroyed.

Operations

Object create(string type)

Create a new object for a given object type. The type is the absolute Slice type id, i.e., the id relative to the unnamed top-level Slice module. For example, the absolute Slice type id for interfaces of type `Bar` in the module `Foo` is `::Foo::Bar`.

The leading `"::"` is required.

Parameters

`type` — The object type.

Return Value

The object created for the given type, or nil if the factory is unable to create the object.

void destroy()

Called when the factory is removed from the communicator, or if the communicator is destroyed.

See Also

- Ice::Communicator::destroy

Copyright 2016, ZeroC, Inc.

Ice-ObjectNotExistException

Ice::ObjectNotExistException

Overview

local exception ObjectNotExistException extends Ice::RequestFailedException

This exception is raised if an object does not exist on the server, that is, if no facets with the given identity exist.

Copyright 2016, ZeroC, Inc.

Ice-ObjectNotFoundException

Ice::ObjectNotFoundException

Overview

exception ObjectNotFoundException

This exception is raised if an object cannot be found.

Used By

- Ice::Locator::findObjectById

Copyright 2016, ZeroC, Inc.

Ice-OpaqueEndpointInfo

Ice::OpaqueEndpointInfo

Overview

local class OpaqueEndpointInfo extends Ice::EndpointInfo

Provides access to the details of an opaque endpoint.

See Also

- Ice::Endpoint

Data Member Index

rawEncoding — The encoding version of the opaque endpoint (to decode or encode the rawBytes).
rawBytes — The raw encoding of the opaque endpoint.

Data Members

Ice::EncodingVersion rawEncoding;

The encoding version of the opaque endpoint (to decode or encode the rawBytes).

Ice::ByteSeq rawBytes;

The raw encoding of the opaque endpoint.

Copyright 2016, ZeroC, Inc.

Ice-OperationInterruptedException

Ice::OperationInterruptedException

Overview

local exception OperationInterruptedException

This exception indicates a request was interrupted.

Copyright 2016, ZeroC, Inc.

Ice-OperationMode

Ice::OperationMode

Overview

enum OperationMode

The OperationMode determines the retry behavior an invocation in case of a (potentially) recoverable error.

Used By

- Current::mode

Enumerator Index

Normal — Ordinary operations have `Normal` mode.
Nonmutating — Operations that use the Slice `nonmutating` keyword must not modify object state.
Idempotent — Operations that use the Slice `idempotent` keyword can modify object state, but invoking an operation twice in a row must result in the same object state as invoking it once.

Enumerators

Normal

Ordinary operations have `Normal` mode. These operations modify object state; invoking such an operation twice in a row has different semantics than invoking it once. The Ice run time guarantees that it will not violate at-most-once semantics for `Normal` operations.

Nonmutating

Operations that use the Slice `nonmutating` keyword must not modify object state. For C++, nonmutating operations generate `const` member functions in the skeleton. In addition, the Ice run time will attempt to transparently recover from certain run-time errors by re-issuing a failed request and propagate the failure to the application only if the second attempt fails.

`Nonmutating` is deprecated; Use the `idempotent` keyword instead. For C++, to retain the mapping of `nonmutating` operations to C++ `const` member functions, use the `["cpp:const"]` metadata directive.

Idempotent

Operations that use the Slice `idempotent` keyword can modify object state, but invoking an operation twice in a row must result in the same object state as invoking it once. For example, `x = 1` is an idempotent statement, whereas `x += 1` is not. For idempotent operations, the Ice run-time uses the same retry behavior as for nonmutating operations in case of a potentially recoverable error.

Copyright 2016, ZeroC, Inc.

Ice-OperationNotExistException

Ice::OperationNotExistException

Overview

local exception OperationNotExistException extends Ice::RequestFailedException

This exception is raised if an operation for a given object does not exist on the server. Typically this is caused by either the client or the server using an outdated Slice specification.

Copyright 2016, ZeroC, Inc.

Ice-Plugin

Ice::Plugin

Overview

local interface Plugin

A communicator plug-in. A plug-in generally adds a feature to a communicator, such as support for a protocol. The communicator loads its plug-ins in two stages: the first stage creates the plug-ins, and the second stage invokes initialize on each one.

Used By

- Ice::PluginManager::addPlugin
- Ice::PluginManager::getPlugin

Operation Index

initialize — Perform any necessary initialization steps.
destroy — Called when the communicator is being destroyed.

Operations

void initialize()

Perform any necessary initialization steps.

void destroy()

Called when the communicator is being destroyed.

Copyright 2016, ZeroC, Inc.

Ice-PluginInitializationException

Ice::PluginInitializationException

Overview

local exception PluginInitializationException

This exception indicates that a failure occurred while initializing a plug-in.

Data Member Index

reason — The reason for the failure.

Data Members

string reason;

The reason for the failure.

Copyright 2016, ZeroC, Inc.

Ice-PluginManager

Ice::PluginManager

Overview

local interface PluginManager

Each communicator has a plug-in manager to administer the set of plug-ins.

Used By

- Ice::Communicator::getPluginManager

Operation Index

initializePlugins — Initialize the configured plug-ins.
getPlugins — Get a list of plugins installed.
getPlugin — Obtain a plug-in by name.
addPlugin — Install a new plug-in.
destroy — Called when the communicator is being destroyed.

Operations

void initializePlugins()

Initialize the configured plug-ins. The communicator automatically initializes the plug-ins by default, but an application may need to interact directly with a plug-in prior to initialization. In this case, the application must set `Ice.InitPlugins=0` and then invoke initializePlugins manually. The plug-ins are initialized in the order in which they are loaded. If a plug-in raises an exception during initialization, the communicator invokes destroy on the plug-ins that have already been initialized.

Exceptions

Ice::InitializationException — Raised if the plug-ins have already been initialized.

Ice::StringSeq getPlugins()

Get a list of plugins installed.

Return Value

The names of the plugins installed.

See Also

- getPlugin

Ice::Plugin getPlugin(string name)

Obtain a plug-in by name.

Parameters

`name` — The plug-in's name.

Return Value

The plug-in.

Exceptions

Ice::NotRegisteredException — Raised if no plug-in is found with the given name.

Copyright 2016, ZeroC, Inc.

void addPlugin(string name, Ice::Plugin pi)

Install a new plug-in.

Parameters

`name` — The plug-in's name.
`pi` — The plug-in.

Exceptions

Ice::AlreadyRegisteredException — Raised if a plug-in already exists with the given name.

void destroy()

Called when the communicator is being destroyed.

Copyright 2016, ZeroC, Inc.

Ice-Process

Ice::Process

Overview

interface Process

An administrative interface for process management. Managed servers must implement this interface.

A servant implementing this interface is a potential target for denial-of-service attacks, therefore proper security precautions should be taken. For example, the servant can use a UUID to make its identity harder to guess, and be registered in an object adapter with a secured endpoint.

Operation Index

shutdown — Initiate a graceful shut-down.
writeMessage — Write a message on the process' stdout or stderr.

Operations

void shutdown()

Initiate a graceful shut-down.

See Also

- Ice::Communicator::shutdown

void writeMessage(string message, int fd)

Write a message on the process' stdout or stderr.

Parameters

`message` — The message.
`fd` — 1 for stdout, 2 for stderr.

Copyright 2016, ZeroC, Inc.

Ice-Properties

Ice::Properties

Overview

local interface Properties

A property set used to configure Ice and Ice applications. Properties are key/value pairs, with both keys and values being strings. By convention, property keys should have the form *application-name*[.*category*[.*sub-category*]].*name*.

Used By

- Ice::Communicator::getProperties
- clone

Operation Index

getProperty — Get a property by key.
getPropertyWithDefault — Get a property by key.
getPropertyAsInt — Get a property as an integer.
getPropertyAsIntWithDefault — Get a property as an integer.
getPropertyAsList — Get a property as a list of strings.
getPropertyAsListWithDefault — Get a property as a list of strings.
getPropertiesForPrefix — Get all properties whose keys begins with *prefix*.
setProperty — Set a property.
getCommandLineOptions — Get a sequence of command-line options that is equivalent to this property set.
parseCommandLineOptions — Convert a sequence of command-line options into properties.
parseIceCommandLineOptions — Convert a sequence of command-line options into properties.
load — Load properties from a file.
clone — Create a copy of this property set.

Operations

string getProperty(string key)

Get a property by key. If the property is not set, an empty string is returned.

Parameters

`key` — The property key.

Return Value

The property value.

See Also

- setProperty

string getPropertyWithDefault(string key, string value)

Get a property by key. If the property is not set, the given default value is returned.

Parameters

`key` — The property key.
`value` — The default value to use if the property does not exist.

Return Value

The property value or the default value.

Copyright 2016, ZeroC, Inc.

See Also

- setProperty

int getPropertyAsInt(string key)

Get a property as an integer. If the property is not set, 0 is returned.

Parameters

`key` — The property key.

Return Value

The property value interpreted as an integer.

See Also

- setProperty

int getPropertyAsIntWithDefault(string key, int value)

Get a property as an integer. If the property is not set, the given default value is returned.

Parameters

`key` — The property key.
`value` — The default value to use if the property does not exist.

Return Value

The property value interpreted as an integer, or the default value.

See Also

- setProperty

Ice::StringSeq getPropertyAsList(string key)

Get a property as a list of strings. The strings must be separated by whitespace or comma. If the property is not set, an empty list is returned. The strings in the list can contain whitespace and commas if they are enclosed in single or double quotes. If quotes are mismatched, an empty list is returned. Within single quotes or double quotes, you can escape the quote in question with
, e.g. O'Reilly can be written as O'Reilly, "O'Reilly" or 'O
'Reilly'.

Parameters

`key` — The property key.

Return Value

The property value interpreted as a list of strings.

See Also

- setProperty

Ice::StringSeq getPropertyAsListWithDefault(string key, Ice::StringSeq value)

Get a property as a list of strings. The strings must be separated by whitespace or comma. If the property is not set, the default list is returned. The strings in the list can contain whitespace and commas if they are enclosed in single or double quotes. If quotes are mismatched, the default list is returned. Within single quotes or double quotes, you can escape the quote in question with
, e.g. O'Reilly can be written as O'Reilly, "O'Reilly" or 'O
'Reilly'.

Parameters

Copyright 2016, ZeroC, Inc.

`key` — The property key.
`value` — The default value to use if the property is not set.

Return Value

The property value interpreted as list of strings, or the default value.

See Also

- setProperty

Ice::PropertyDict getPropertiesForPrefix(string prefix)

Get all properties whose keys begins with *prefix*. If *prefix* is an empty string, then all properties are returned.

Parameters

`prefix` — The prefix to search for (empty string if none).

Return Value

The matching property set.

void setProperty(string key, string value)

Set a property. To unset a property, set it to the empty string.

Parameters

`key` — The property key.
`value` — The property value.

See Also

- getProperty

Ice::StringSeq getCommandLineOptions()

Get a sequence of command-line options that is equivalent to this property set. Each element of the returned sequence is a command-line option of the form `--key=value`.

Return Value

The command line options for this property set.

Ice::StringSeq parseCommandLineOptions(string prefix, Ice::StringSeq options)

Convert a sequence of command-line options into properties. All options that begin with `--prefix.` are converted into properties. If the prefix is empty, all options that begin with `--` are converted to properties.

Parameters

`prefix` — The property prefix, or an empty string to convert all options starting with `--`.
`options` — The command-line options.

Return Value

The command-line options that do not start with the specified prefix, in their original order.

Ice::StringSeq parseIceCommandLineOptions(Ice::StringSeq options)

Convert a sequence of command-line options into properties. All options that begin with one of the following prefixes are converted into properties: `--Ice`, `--IceBox`, `--IceGrid`, `--IcePatch2`, `--IceSSL`, `--IceStorm`, `--Freeze`, and `--Glacier2`.

Copyright 2016, ZeroC, Inc.

Parameters

`options` — The command-line options.

Return Value

The command-line options that do not start with one of the listed prefixes, in their original order.

void load(string file)

Load properties from a file.

Parameters

`file` — The property file.

Ice::Properties clone()

Create a copy of this property set.

Return Value

A copy of this property set.

Copyright 2016, ZeroC, Inc.

Ice-PropertiesAdmin

Ice::PropertiesAdmin

Overview

interface PropertiesAdmin

The PropertiesAdmin interface provides remote access to the properties of a communicator.

Operation Index

getProperty — Get a property by key.
getPropertiesForPrefix — Get all properties whose keys begin with *prefix*.
setProperties — Update the communicator's properties with the given property set.

Operations

string getProperty(string key)

Get a property by key. If the property is not set, an empty string is returned.

Parameters

`key` — The property key.

Return Value

The property value.

["java:type:java.util.TreeMap<String, String>"] Ice::PropertyDict getPropertiesForPrefix(string prefix)

Get all properties whose keys begin with *prefix*. If *prefix* is an empty string then all properties are returned.

Parameters

`prefix` — The prefix to search for (empty string if none).

Return Value

The matching property set.

["amd"] void setProperties(Ice::PropertyDict newProperties)

Update the communicator's properties with the given property set.

Parameters

`newProperties` — Properties to be added, changed, or removed. If an entry in *newProperties* matches the name of an existing property, that property's value is replaced with the new value. If the new value is an empty string, the property is removed. Any existing properties that are not modified or removed by the entries in newProperties are retained with their original values.

Copyright 2016, ZeroC, Inc.

Ice-ProtocolException

Ice::ProtocolException

Overview

local exception ProtocolException

A generic exception base for all kinds of protocol error conditions.

Derived Exceptions

- Ice::BadMagicException
- Ice::CloseConnectionException
- Ice::CompressionException
- Ice::ConnectionNotValidatedException
- Ice::DatagramLimitException
- Ice::ForcedCloseConnectionException
- Ice::IllegalMessageSizeException
- Ice::MarshalException
- Ice::UnknownMessageException
- Ice::UnknownReplyStatusException
- Ice::UnknownRequestIdException
- Ice::UnsupportedEncodingException
- Ice::UnsupportedProtocolException

Data Member Index

reason — The reason for the failure.

Data Members

string reason;

The reason for the failure.

Copyright 2016, ZeroC, Inc.

Ice-ProtocolVersion

Ice::ProtocolVersion

Overview

struct ProtocolVersion

A version structure for the protocol version.

Used By

- Ice::UnsupportedProtocolException::bad
- Ice::UnsupportedProtocolException::supported

Data Member Index

major
minor

Data Members

byte major;

byte minor;

Copyright 2016, ZeroC, Inc.

Ice-ProxyParseException

Ice::ProxyParseException

Overview

local exception ProxyParseException

This exception is raised if there was an error while parsing a stringified proxy.

Data Member Index

str — Describes the failure and includes the string that could not be parsed.

Data Members

string str;

Describes the failure and includes the string that could not be parsed.

Copyright 2016, ZeroC, Inc.

Ice-ProxyUnmarshalException

Ice::ProxyUnmarshalException

Overview

local exception ProxyUnmarshalException extends Ice::MarshalException

This exception is raised if inconsistent data is received while unmarshaling a proxy.

Copyright 2016, ZeroC, Inc.

Ice-RemoteLogger

Ice::RemoteLogger

Overview

interface RemoteLogger

The Ice remote logger interface. An application can implement a RemoteLogger to receive the log messages sent to the local Ice::Logger of another Ice application.

Operation Index

init — init is called by Ice::LoggerAdmin::attachRemoteLogger when a RemoteLogger proxy is attached.
log — Log a Ice::LogMessage.

Operations

void init(string prefix, **Ice::LogMessageSeq** logMessages)

init is called by Ice::LoggerAdmin::attachRemoteLogger when a RemoteLogger proxy is attached.

Parameters

`prefix` — The prefix of the associated local Ice::Logger
`logMessages` — Old log messages generated before "now".

See Also

- Ice::LoggerAdmin::attachRemoteLogger

void log(**Ice::LogMessage** message)

Log a Ice::LogMessage. Note that log may be called by Ice::LoggerAdmin before init.

Parameters

`message` — The message to log.

See Also

- Ice::Logger

Copyright 2016, ZeroC, Inc.

Ice-RemoteLoggerAlreadyAttachedException

Ice::RemoteLoggerAlreadyAttachedException

Overview

exception RemoteLoggerAlreadyAttachedException

An exception thrown by Ice::LoggerAdmin::attachRemoteLogger to report that the provided Ice::RemoteLogger was previously attached to this Ice::LoggerAdmin.

Used By

- Ice::LoggerAdmin::attachRemoteLogger

Copyright 2016, ZeroC, Inc.

Ice-RequestFailedException

Ice::RequestFailedException

Overview

local exception RequestFailedException

This exception is raised if a request failed. This exception, and all exceptions derived from Ice::RequestFailedException, are transmitted by the Ice protocol, even though they are declared `local`.

Derived Exceptions

- Ice::FacetNotExistException
- Ice::ObjectNotExistException
- Ice::OperationNotExistException

Data Member Index

id — The identity of the Ice Object to which the request was sent.
facet — The facet to which the request was sent.
operation — The operation name of the request.

Data Members

Ice::Identity id;

The identity of the Ice Object to which the request was sent.

string facet;

The facet to which the request was sent.

string operation;

The operation name of the request.

Copyright 2016, ZeroC, Inc.

Ice-ResponseSentException

Ice::ResponseSentException

Overview

local exception ResponseSentException

Indicates that the response to a request has already been sent; re-dispatching such a request is not possible.

Copyright 2016, ZeroC, Inc.

Ice-Router

Ice::Router

Overview

interface Router

The Ice router interface. Routers can be set either globally with Ice::Communicator::setDefaultRouter, or with `ice_router` on specific proxies.

Derived Classes and Interfaces

- Glacier2::Router

Operation Index

getClientProxy — Get the router's client proxy, i.e., the proxy to use for forwarding requests from the client to the router.
getServerProxy — Get the router's server proxy, i.e., the proxy to use for forwarding requests from the server to the router.
addProxies — Add new proxy information to the router's routing table.

Operations

["nonmutating"] Object* getClientProxy()

Get the router's client proxy, i.e., the proxy to use for forwarding requests from the client to the router. If a null proxy is returned, the client will forward requests to the router's endpoints.

Return Value

The router's client proxy.

["nonmutating"] Object* getServerProxy()

Get the router's server proxy, i.e., the proxy to use for forwarding requests from the server to the router.

Return Value

The router's server proxy.

Ice::ObjectProxySeq addProxies(Ice::ObjectProxySeq proxies)

Add new proxy information to the router's routing table.

Parameters

`proxies` — The proxies to add.

Return Value

Proxies discarded by the router.

Copyright 2016, ZeroC, Inc.

Ice-RouterFinder

Ice::RouterFinder

Overview

interface RouterFinder

This inferface should be implemented by services implementing the Ice::Router interface. It should be advertised through an Ice object with the identity `Ice/RouterFinder'. This allows clients to retrieve the router proxy with just the endpoint information of the service.

Operation Index

getRouter — Get the router proxy implemented by the process hosting this finder object.

Operations

Ice::Router* getRouter()

Get the router proxy implemented by the process hosting this finder object. The proxy might point to several replicas.

Return Value

The router proxy.

Copyright 2016, ZeroC, Inc.

Ice-SecurityException

Ice::SecurityException

Overview

local exception SecurityException

This exception indicates a failure in a security subsystem, such as the IceSSL plug-in.

Data Member Index

reason — The reason for the failure.

Data Members

string reason;

The reason for the failure.

Copyright 2016, ZeroC, Inc.

Ice-ServantLocator

Ice::ServantLocator

Overview

local interface ServantLocator

A servant locator is called by an object adapter to locate a servant that is not found in its active servant map.

Derived Classes and Interfaces

- Freeze::Evictor

Used By

- Ice::ObjectAdapter::addServantLocator
- Ice::ObjectAdapter::findServantLocator
- Ice::ObjectAdapter::removeServantLocator

See Also

- Ice::ObjectAdapter
- Ice::ObjectAdapter::addServantLocator
- Ice::ObjectAdapter::findServantLocator

Operation Index

locate — Called before a request is dispatched if a servant cannot be found in the object adapter's active servant map.
finished — Called by the object adapter after a request has been made.
deactivate — Called when the object adapter in which this servant locator is installed is destroyed.

Operations

Object locate(**Ice::Current** curr, out LocalObject cookie)

Called before a request is dispatched if a servant cannot be found in the object adapter's active servant map. Note that the object adapter does not automatically insert the returned servant into its active servant map. This must be done by the servant locator implementation, if this is desired. `locate` can throw any user exception. If it does, that exception is marshaled back to the client. If the Slice definition for the corresponding operation includes that user exception, the client receives that user exception; otherwise, the client receives Ice::UnknownUserException. If `locate` throws any exception, the Ice run time does *not* call `finished`.

If you call `locate` from your own code, you must also call `finished` when you have finished using the servant, provided that `locate` returned a non-null servant; otherwise, you will get undefined behavior if you use servant locators such as the Freeze Evictor.

Parameters

`curr` — Information about the current operation for which a servant is required.
`cookie` — A "cookie" that will be passed to `finished`.

Return Value

The located servant, or null if no suitable servant has been found.

See Also

- Ice::ObjectAdapter
- Ice::Current
- finished

void finished(**Ice::Current** curr, Object servant, LocalObject cookie)

Called by the object adapter after a request has been made. This operation is only called if `locate` was called prior to the request and

Copyright 2016, ZeroC, Inc.

returned a non-null servant. This operation can be used for cleanup purposes after a request. `finished` can throw any user exception. If it does, that exception is marshaled back to the client. If the Slice definition for the corresponding operation includes that user exception, the client receives that user exception; otherwise, the client receives Ice::UnknownUserException. If both the operation and `finished` throw an exception, the exception thrown by `finished` is marshaled back to the client.

Parameters

`curr` — Information about the current operation call for which a servant was located by `locate`.
`servant` — The servant that was returned by `locate`.
`cookie` — The cookie that was returned by `locate`.

See Also

- Ice::ObjectAdapter
- Ice::Current
- locate

void deactivate(string category)

Called when the object adapter in which this servant locator is installed is destroyed.

Parameters

`category` — Indicates for which category the servant locator is being deactivated.

See Also

- Ice::ObjectAdapter::destroy
- Ice::Communicator::shutdown
- Ice::Communicator::destroy

Copyright 2016, ZeroC, Inc.

Ice-ServerNotFoundException

Ice::ServerNotFoundException

Overview

exception ServerNotFoundException

This exception is raised if a server cannot be found.

Used By

- Ice::LocatorRegistry::setServerProcessProxy

Copyright 2016, ZeroC, Inc.

Ice-SocketException

Ice::SocketException

Overview

local exception SocketException extends Ice::SyscallException

This exception indicates socket errors.

Derived Exceptions

- Ice::ConnectFailedException
- Ice::ConnectionLostException

Copyright 2016, ZeroC, Inc.

Ice-StringConversionException

Ice::StringConversionException

Overview

local exception StringConversionException extends Ice::MarshalException

This exception is raised when a string conversion to or from UTF-8 fails during marshaling or unmarshaling.

Copyright 2016, ZeroC, Inc.

Ice-SyscallException

Ice::SyscallException

Overview

local exception SyscallException

This exception is raised if a system error occurred in the server or client process. There are many possible causes for such a system exception. For details on the cause, error should be inspected.

Derived Exceptions

- Ice::FileException
- Ice::SocketException

Data Member Index

error — The error number describing the system exception.

Data Members

int error;

The error number describing the system exception. For C++ and Unix, this is equivalent to `errno`. For C++ and Windows, this is the value returned by `GetLastError()` or `WSAGetLastError()`.

Copyright 2016, ZeroC, Inc.

Ice-TCPConnectionInfo

Ice::TCPConnectionInfo

Overview

local class TCPConnectionInfo extends Ice::IPConnectionInfo

Provides access to the connection details of a TCP connection

Derived Classes and Interfaces

- Ice::WSConnectionInfo

Copyright 2016, ZeroC, Inc.

Ice-TCPEndpointInfo

Ice::TCPEndpointInfo

Overview

local class TCPEndpointInfo extends Ice::IPEndpointInfo

Provides access to a TCP endpoint information.

Derived Classes and Interfaces

- Ice::WSEndpointInfo

See Also

- Ice::Endpoint

Copyright 2016, ZeroC, Inc.

Ice-TimeoutException

Ice::TimeoutException

Overview

local exception TimeoutException

This exception indicates a timeout condition.

Derived Exceptions

- Ice::CloseTimeoutException
- Ice::ConnectTimeoutException
- Ice::ConnectionTimeoutException
- Ice::InvocationTimeoutException

Copyright 2016, ZeroC, Inc.

Ice-TwowayOnlyException

Ice::TwowayOnlyException

Overview

local exception TwowayOnlyException

The operation can only be invoked with a twoway request. This exception is raised if an attempt is made to invoke an operation with `ice_oneway`, `ice_batchOneway`, `ice_datagram`, or `ice_batchDatagram` and the operation has a return value, out-parameters, or an exception specification.

Data Member Index

operation — The name of the operation that was invoked.

Data Members

string operation;

The name of the operation that was invoked.

Copyright 2016, ZeroC, Inc.

Ice-UDPConnectionInfo

Ice::UDPConnectionInfo

Overview

local class UDPConnectionInfo extends Ice::IPConnectionInfo

Provides access to the connection details of a UDP connection

Data Member Index

mcastAddress — The multicast address.
mcastPort — The multicast port.

Data Members

string mcastAddress;

The multicast address.

int mcastPort;

The multicast port.

Copyright 2016, ZeroC, Inc.

Ice-UDPEndpointInfo

Ice::UDPEndpointInfo

Overview

local class UDPEndpointInfo extends Ice::IPEndpointInfo

Provides access to an UDP endpoint information.

See Also

- Ice::Endpoint

Data Member Index

mcastInterface — The multicast interface.
mcastTtl — The multicast time-to-live (or hops).

Data Members

string mcastInterface;

The multicast interface.

int mcastTtl;

The multicast time-to-live (or hops).

Copyright 2016, ZeroC, Inc.

Ice-UnexpectedObjectException

Ice::UnexpectedObjectException

Overview

local exception UnexpectedObjectException extends Ice::MarshalException

This exception is raised if the type of an unmarshaled Slice class instance does not match its expected type. This can happen if client and server are compiled with mismatched Slice definitions or if a class of the wrong type is passed as a parameter or return value using dynamic invocation. This exception can also be raised if IceStorm is used to send Slice class instances and an operation is subscribed to the wrong topic.

Data Member Index

type — The Slice type ID of the class instance that was unmarshaled.
expectedType — The Slice type ID that was expected by the receiving operation.

Data Members

string type;

The Slice type ID of the class instance that was unmarshaled.

string expectedType;

The Slice type ID that was expected by the receiving operation.

Copyright 2016, ZeroC, Inc.

Ice-UnknownException

Ice::UnknownException

Overview

local exception UnknownException

This exception is raised if an operation call on a server raises an unknown exception. For example, for C++, this exception is raised if the server throws a C++ exception that is not directly or indirectly derived from `Ice::LocalException` or `Ice::UserException`.

Derived Exceptions

- Ice::UnknownLocalException
- Ice::UnknownUserException

Data Member Index

unknown — This field is set to the textual representation of the unknown exception if available.

Data Members

string unknown;

This field is set to the textual representation of the unknown exception if available.

Copyright 2016, ZeroC, Inc.

Ice-UnknownLocalException

Ice::UnknownLocalException

Overview

local exception UnknownLocalException extends Ice::UnknownException

This exception is raised if an operation call on a server raises a local exception. Because local exceptions are not transmitted by the Ice protocol, the client receives all local exceptions raised by the server as Ice::UnknownLocalException. The only exception to this rule are all exceptions derived from Ice::RequestFailedException, which are transmitted by the Ice protocol even though they are declared `local`.

Copyright 2016, ZeroC, Inc.

Ice-UnknownMessageException

Ice::UnknownMessageException

Overview

local exception UnknownMessageException extends Ice::ProtocolException

This exception indicates that an unknown protocol message has been received.

Copyright 2016, ZeroC, Inc.

Ice-UnknownReplyStatusException

Ice::UnknownReplyStatusException

Overview

local exception UnknownReplyStatusException extends Ice::ProtocolException

This exception indicates that an unknown reply status has been received.

Copyright 2016, ZeroC, Inc.

Ice-UnknownRequestIdException

Ice::UnknownRequestIdException

Overview

local exception UnknownRequestIdException extends Ice::ProtocolException

This exception indicates that a response for an unknown request ID has been received.

Copyright 2016, ZeroC, Inc.

Ice-UnknownUserException

Ice::UnknownUserException

Overview

local exception UnknownUserException extends Ice::UnknownException

An operation raised an incorrect user exception. This exception is raised if an operation raises a user exception that is not declared in the exception's `throws` clause. Such undeclared exceptions are not transmitted from the server to the client by the Ice protocol, but instead the client just gets an Ice::UnknownUserException. This is necessary in order to not violate the contract established by an operation's signature: Only local exceptions and user exceptions declared in the `throws` clause can be raised.

Copyright 2016, ZeroC, Inc.

Ice-UnmarshalOutOfBoundsException

Ice::UnmarshalOutOfBoundsException

Overview

local exception UnmarshalOutOfBoundsException extends Ice::MarshalException

This exception is raised if an out-of-bounds condition occurs during unmarshaling.

Copyright 2016, ZeroC, Inc.

Ice-UnsupportedEncodingException

Ice::UnsupportedEncodingException

Overview

local exception UnsupportedEncodingException extends Ice::ProtocolException

This exception indicates an unsupported data encoding version.

Data Member Index

bad — The version of the unsupported encoding.
supported — The version of the encoding that is supported.

Data Members

Ice::EncodingVersion bad;

The version of the unsupported encoding.

Ice::EncodingVersion supported;

The version of the encoding that is supported.

Copyright 2016, ZeroC, Inc.

Ice-UnsupportedProtocolException

Ice::UnsupportedProtocolException

Overview

local exception UnsupportedProtocolException extends Ice::ProtocolException

This exception indicates an unsupported protocol version.

Data Member Index

bad — The version of the unsupported protocol.
supported — The version of the protocol that is supported.

Data Members

Ice::ProtocolVersion bad;

The version of the unsupported protocol.

Ice::ProtocolVersion supported;

The version of the protocol that is supported.

Copyright 2016, ZeroC, Inc.

Ice-VersionMismatchException

Ice::VersionMismatchException

Overview

local exception VersionMismatchException

This exception is raised if the Ice library version does not match the version in the Ice header files.

Copyright 2016, ZeroC, Inc.

Ice-VersionParseException

Ice::VersionParseException

Overview

local exception VersionParseException

This exception is raised if there was an error while parsing a version.

Data Member Index

str — Describes the failure and includes the string that could not be parsed.

Data Members

string str;

Describes the failure and includes the string that could not be parsed.

Copyright 2016, ZeroC, Inc.

Ice-WSConnectionInfo

Ice::WSConnectionInfo

Overview

local class WSConnectionInfo extends Ice::TCPConnectionInfo

Provides access to the connection details of a WebSocket connection

Data Member Index

headers — The headers from the HTTP upgrade request.

Data Members

Ice::HeaderDict headers;

The headers from the HTTP upgrade request.

Copyright 2016, ZeroC, Inc.

Ice-WSEndpointInfo

Ice::WSEndpointInfo

Overview

local class WSEndpointInfo extends Ice::TCPEndpointInfo

Provides access to a WebSocket endpoint information.

Data Member Index

resource — The URI configured with the endpoint.

Data Members

string resource;

The URI configured with the endpoint.

Copyright 2016, ZeroC, Inc.

IceSSL Slice API

IceSSL

Overview

["objc:prefix:ICESSL"] module IceSSL

IceSSL provides a secure transport for Ice.

Class Index

ConnectionInfo — Provides access to the connection details of an SSL connection
EndpointInfo — Provides access to an SSL endpoint information.
WSSConnectionInfo — Provides access to the connection details of a secure WebSocket connection
WSSEndpointInfo — Provides access to a secure WebSocket endpoint information.

Constant Index

EndpointType — Uniquely identifies SSL endpoints.

Constants

const short EndpointType = 2;

Uniquely identifies SSL endpoints.

Copyright 2016, ZeroC, Inc.

IceSSL-ConnectionInfo

IceSSL::ConnectionInfo

Overview

local class ConnectionInfo extends Ice::IPConnectionInfo

Provides access to the connection details of an SSL connection

Derived Classes and Interfaces

- IceSSL::WSSConnectionInfo

Data Member Index

cipher — The negotiated cipher suite.
certs — The certificate chain.
verified — The certificate chain verification status.

Data Members

string cipher;

The negotiated cipher suite.

Ice::StringSeq certs;

The certificate chain.

bool verified;

The certificate chain verification status.

Copyright 2016, ZeroC, Inc.

IceSSL-EndpointInfo

IceSSL::EndpointInfo

Overview

local class EndpointInfo extends Ice::IPEndpointInfo

Provides access to an SSL endpoint information.

Derived Classes and Interfaces

- IceSSL::WSSEndpointInfo

Copyright 2016, ZeroC, Inc.

IceSSL-WSSConnectionInfo

IceSSL::WSSConnectionInfo

Overview

local class WSSConnectionInfo extends IceSSL::ConnectionInfo

Provides access to the connection details of a secure WebSocket connection

Data Member Index

headers — The headers from the HTTP upgrade request.

Data Members

Ice::HeaderDict headers;

The headers from the HTTP upgrade request.

Copyright 2016, ZeroC, Inc.

IceSSL-WSSEndpointInfo

IceSSL::WSSEndpointInfo

Overview

local class WSSEndpointInfo extends IceSSL::EndpointInfo

Provides access to a secure WebSocket endpoint information.

Data Member Index

resource — The URI configured with the endpoint.

Data Members

string resource;

The URI configured with the endpoint.

Copyright 2016, ZeroC, Inc.

IceStorm Slice API

IceStorm

Overview

["objc:prefix:ICESTORM"] module IceStorm

A messaging service with support for federation. In contrast to most other messaging or event services, IceStorm supports typed events, meaning that broadcasting a message over a federation is as easy as invoking a method on an interface.

Interface Index

Finder — This inferface is advertised by the IceStorm service through the Ice object with the identity `IceStorm/Finder'.
Topic — Publishers publish information on a particular topic.
TopicManager — A topic manager manages topics, and subscribers to topics.

Exception Index

AlreadySubscribed — This exception indicates that an attempt was made to subscribe a proxy for which a subscription already exists.
BadQoS — This exception indicates that a subscription failed due to an invalid QoS.
InvalidSubscriber — This exception indicates that an attempt was made to subscribe a proxy that is null.
LinkExists — This exception indicates that an attempt was made to create a link that already exists.
NoSuchLink — This exception indicates that an attempt was made to remove a link that does not exist.
NoSuchTopic — This exception indicates that an attempt was made to retrieve a topic that does not exist.
TopicExists — This exception indicates that an attempt was made to create a topic that already exists.

Structure Index

LinkInfo — Information on the topic links.

Sequence Index

LinkInfoSeq — A sequence of LinkInfo objects.

Dictionary Index

QoS — This dictionary represents quality of service parameters.
TopicDict — Mapping of topic name to topic proxy.

Sequences

sequence<LinkInfo> LinkInfoSeq

A sequence of LinkInfo objects.

Used By

- Topic::getLinkInfoSeq

Dictionaries

dictionary<string, string> QoS

This dictionary represents quality of service parameters.

Copyright 2016, ZeroC, Inc.

Used By

- Topic::subscribeAndGetPublisher

See Also

- Topic::subscribeAndGetPublisher

dictionary<string, Topic*> TopicDict

Mapping of topic name to topic proxy.

Used By

- TopicManager::retrieveAll

Copyright 2016, ZeroC, Inc.

IceStorm-AlreadySubscribed

IceStorm::AlreadySubscribed

Overview

exception AlreadySubscribed

This exception indicates that an attempt was made to subscribe a proxy for which a subscription already exists.

Used By

- IceStorm::Topic::subscribeAndGetPublisher

Copyright 2016, ZeroC, Inc.

IceStorm-BadQoS

IceStorm::BadQoS

Overview

exception BadQoS

This exception indicates that a subscription failed due to an invalid QoS.

Used By

- IceStorm::Topic::subscribeAndGetPublisher

Data Member Index

reason — The reason for the failure.

Data Members

string reason;

The reason for the failure.

Copyright 2016, ZeroC, Inc.

IceStorm-Finder

IceStorm::Finder

Overview

interface Finder

This inferface is advertised by the IceStorm service through the Ice object with the identity `IceStorm/Finder'. This allows clients to retrieve the topic manager with just the endpoint information of the IceStorm service.

Operation Index

getTopicManager — Get the topic manager proxy.

Operations

IceStorm::TopicManager* getTopicManager()

Get the topic manager proxy. The proxy might point to several replicas.

Return Value

The topic manager proxy.

Copyright 2016, ZeroC, Inc.

IceStorm-InvalidSubscriber

IceStorm::InvalidSubscriber

Overview

exception InvalidSubscriber

This exception indicates that an attempt was made to subscribe a proxy that is null.

Used By

- IceStorm::Topic::subscribeAndGetPublisher

Data Member Index

reason — The reason for the failure.

Data Members

string reason;

The reason for the failure.

Copyright 2016, ZeroC, Inc.

IceStorm-LinkExists

IceStorm::LinkExists

Overview

exception LinkExists

This exception indicates that an attempt was made to create a link that already exists.

Used By

- IceStorm::Topic::link

Data Member Index

name — The name of the linked topic.

Data Members

string name;

The name of the linked topic.

Copyright 2016, ZeroC, Inc.

IceStorm-LinkInfo

IceStorm::LinkInfo

Overview

struct LinkInfo

Information on the topic links.

Used By

- IceStorm::LinkInfoSeq

Data Member Index

theTopic — The linked topic.
name — The name of the linked topic.
cost — The cost of traversing this link.

Data Members

IceStorm::Topic* theTopic;

The linked topic.

string name;

The name of the linked topic.

int cost;

The cost of traversing this link.

Copyright 2016, ZeroC, Inc.

IceStorm-NoSuchLink

IceStorm::NoSuchLink

Overview

exception NoSuchLink

This exception indicates that an attempt was made to remove a link that does not exist.

Used By

- IceStorm::Topic::unlink

Data Member Index

name — The name of the link that does not exist.

Data Members

string name;

The name of the link that does not exist.

Copyright 2016, ZeroC, Inc.

IceStorm-NoSuchTopic

IceStorm::NoSuchTopic

Overview

exception NoSuchTopic

This exception indicates that an attempt was made to retrieve a topic that does not exist.

Used By

- IceStorm::TopicManager::retrieve

Data Member Index

name — The name of the topic that does not exist.

Data Members

string name;

The name of the topic that does not exist.

Copyright 2016, ZeroC, Inc.

IceStorm-Topic

IceStorm::Topic

Overview

interface Topic

Publishers publish information on a particular topic. A topic logically represents a type.

See Also

- IceStorm::TopicManager

Operation Index

getName — Get the name of this topic.
getPublisher — Get a proxy to a publisher object for this topic.
getNonReplicatedPublisher — Get a non-replicated proxy to a publisher object for this topic.
subscribeAndGetPublisher — Subscribe with the given `qos` to this topic.
unsubscribe — Unsubscribe the given `subscriber`.
link — Create a link to the given topic.
unlink — Destroy the link from this topic to the given topic `linkTo`.
getLinkInfoSeq — Retrieve information on the current links.
getSubscribers — Retrieve the list of subscribers for this topic.
destroy — Destroy the topic.

Operations

["nonmutating"] string getName()

Get the name of this topic.

Return Value

The name of the topic.

See Also

- IceStorm::TopicManager::create

["nonmutating"] Object* getPublisher()

Get a proxy to a publisher object for this topic. To publish data to a topic, the publisher calls getPublisher and then casts to the topic type. An unchecked cast must be used on this proxy. If a replicated IceStorm deployment is used this call may return a replicated proxy.

Return Value

A proxy to publish data on this topic.

["nonmutating"] Object* getNonReplicatedPublisher()

Get a non-replicated proxy to a publisher object for this topic. To publish data to a topic, the publisher calls getPublisher and then casts to the topic type. An unchecked cast must be used on this proxy.

Return Value

A proxy to publish data on this topic.

Object* subscribeAndGetPublisher(IceStorm::QoS theQoS, Object* subscriber) throws IceStorm::AlreadySubscribed, IceStorm::InvalidSubscriber, IceStorm::BadQoS

Copyright 2016, ZeroC, Inc.

Subscribe with the given `qos` to this topic. A per-subscriber publisher object is returned.

Parameters

`theQoS` — The quality of service parameters for this subscription.
`subscriber` — The subscriber's proxy.

Return Value

The per-subscriber publisher object.

Exceptions

IceStorm::AlreadySubscribed — Raised if the subscriber object is already subscribed.
`NullSubscriber` — Raised if the subscriber object is null.
IceStorm::BadQoS — Raised if the requested quality of service is unavailable or invalid.

See Also

- unsubscribe

void unsubscribe(Object* subscriber)

Unsubscribe the given `subscriber`.

Parameters

`subscriber` — The proxy of an existing subscriber.

See Also

- subscribeAndGetPublisher

void link(IceStorm::Topic* linkTo, int cost) throws IceStorm::LinkExists

Create a link to the given topic. All events originating on this topic will also be sent to `linkTo`.

Parameters

`linkTo` — The topic to link to.
`cost` — The cost to the linked topic.

Exceptions

IceStorm::LinkExists — Raised if a link to the same topic already exists.

void unlink(IceStorm::Topic* linkTo) throws IceStorm::NoSuchLink

Destroy the link from this topic to the given topic `linkTo`.

Parameters

`linkTo` — The topic to destroy the link to.

Exceptions

IceStorm::NoSuchLink — Raised if a link to the topic does not exist.

["nonmutating"] IceStorm::LinkInfoSeq getLinkInfoSeq()

Retrieve information on the current links.

Return Value

A sequence of LinkInfo objects.

Copyright 2016, ZeroC, Inc.

["nonmutating"] Ice::IdentitySeq getSubscribers()

Retrieve the list of subscribers for this topic.

Return Value

The sequence of Ice identities for the subscriber objects.

void destroy()

Destroy the topic.

Copyright 2016, ZeroC, Inc.

IceStorm-TopicExists

IceStorm::TopicExists

Overview

exception TopicExists

This exception indicates that an attempt was made to create a topic that already exists.

Used By

- IceStorm::TopicManager::create

Data Member Index

name — The name of the topic that already exists.

Data Members

string name;

The name of the topic that already exists.

Copyright 2016, ZeroC, Inc.

IceStorm-TopicManager

IceStorm::TopicManager

Overview

interface TopicManager

A topic manager manages topics, and subscribers to topics.

See Also

- IceStorm::Topic

Operation Index

create — Create a new topic.
retrieve — Retrieve a topic by name.
retrieveAll — Retrieve all topics managed by this topic manager.
getSliceChecksums — Returns the checksums for the IceStorm Slice definitions.

Operations

IceStorm::Topic* create(string name) throws IceStorm::TopicExists

Create a new topic. The topic name must be unique, otherwise IceStorm::TopicExists is raised.

Parameters

`name` — The name of the topic.

Return Value

A proxy to the topic instance.

Exceptions

IceStorm::TopicExists — Raised if a topic with the same name already exists.

["nonmutating"] IceStorm::Topic* retrieve(string name) throws IceStorm::NoSuchTopic

Retrieve a topic by name.

Parameters

`name` — The name of the topic.

Return Value

A proxy to the topic instance.

Exceptions

IceStorm::NoSuchTopic — Raised if the topic does not exist.

["nonmutating"] IceStorm::TopicDict retrieveAll()

Retrieve all topics managed by this topic manager.

Return Value

A dictionary of string, topic proxy pairs.

Copyright 2016, ZeroC, Inc.

["nonmutating"] Ice::SliceChecksumDict getSliceChecksums()

Returns the checksums for the IceStorm Slice definitions.

Return Value

A dictionary mapping Slice type ids to their checksums.

Copyright 2016, ZeroC, Inc.