

Java API User Guide

Release 6.1

February 2003

Java API User Guide

ObjectStore Release 6.1 for all platforms, February 2003

© 2003 Progress Software Corporation. All rights reserved.

Progress® software products are copyrighted and all rights are reserved by Progress Software Corporation. This manual is also copyrighted and all rights are reserved. This manual may not, in whole or in part, be copied, photocopied, translated, or reduced to any electronic medium or machine-readable form without prior consent, in writing, from Progress Software Corporation.

The information in this manual is subject to change without notice, and Progress Software Corporation assumes no responsibility for any errors that may appear in this document.

The references in this manual to specific platforms supported are subject to change.

Allegrix, Leadership by Design, Object Design, ObjectStore, Progress, Powered by Progress, Progress Fast Track, Progress Profiles, Partners in Progress, Partners en Progress, Progress en Partners, Progress in Progress, P.I.P., Progress Results, ProVision, ProCare, ProtoSpeed, SmartBeans, SpeedScript, and WebSpeed are registered trademarks of Progress Software Corporation or one of its subsidiaries or affiliates in the U.S. and other countries. A Data Center of Your Very Own, Apptivity, AppsAlive, AppServer, ASPen, ASP-in-a-Box, BPM, Cache-Forward, Empowerment Center, eXcelon, EXLN, Fathom, Future Proof, Progress for Partners, IntelliStream, Javlin, ObjectStore Browsers, OpenEdge, POSSE, POSSENET, Progress Dynamics, Progress Software Developers Network, RTEE, Schemadesigner, SectorAlliance, SmartBrowser, SmartComponent, SmartDataBrowser, SmartDataObjects, SmartDataView, SmartDialog, SmartFolder, SmartFrame, SmartObjects, SmartPanel, SmartQuery, SmartViewer, SmartWindow, Stylus, Stylus Studio, WebClient, Who Makes Progress, XIS, XIS Lite, and XPress are trademarks or service marks of Progress Software Corporation or one of its subsidiaries or affiliates in the U.S. and other countries.

Any other trademarks and service marks contained herein may be the property of their respective owners.

Contents

Preface	xix
Introducing ObjectStore	1
What Is ObjectStore?	1
What ObjectStore Does	2
Benefits of Using ObjectStore	3
Description of ObjectStore Process Architecture.	3
Definitions of ObjectStore Terms	5
Session	5
Persistence Capable	6
Persistent Object	6
Persistence Aware.	9
Primary Object	10
Peer Object	10
Transient Object.	10
Transitive Persistence	11
Annotations	11
Database Roots	11
Prerequisites for Using the ObjectStore Java Interface	11
Example of Using ObjectStore	13
Overview of Required Components	13
Sample Code.	14
Before You Run the Program	17
Adding an Entry to CLASSPATH	17
Compiling the Program	17

Running the Postprocessor	18
Running the Program	18
Using Sessions to Manage Threads	19
How Sessions Keep Threads Organized	19
What Is a Session?	20
How Are Threads Related to Sessions?	20
What Is the Benefit of a Session?	21
What Kinds of Sessions Are There?	22
Creating Sessions	23
Creating Global Sessions	23
Creating Nonglobal Sessions	24
Working with Sessions.	25
Sessions and Transactions	25
Shutting Down Sessions.	26
Obtaining a Session	26
Determining Whether a Session Is Active	27
Associating Threads with Sessions	27
Joining Threads to a Session Automatically.	27
Associating a Persistent Object with a Session	28
Rules for Joining a Thread to a Session Automatically	29
Examples of Calls That Imply Sessions.	29
Examples of Calls That Do Not Imply Sessions	29
Explicitly Associating Threads with a Session	30
Working with Threads	31
Cooperating Threads	31
Noncooperating Threads.	32
Synchronizing Threads	32
Removing Threads from Sessions	32
Threads That Create a Session	33
Other Threads.	33
Determining Whether ObjectStore Is Initialized for the Current Thread	33
Threads and Persistent Objects.	34
Multiple Representations of the Same Object	34

Example of Multiple Sessions	35
Application Responsibility.	35
Effects of Committing a Transaction	35
API Objects and Sessions.	36
Description of Concurrency Rules.	36
Granularity of Concurrency	36
Converting Read Locks to Write Locks.	37
Description of ObjectStore Properties.	37
About Property Lists Relevant to ObjectStore	37
Description of com.odi.addressSpaceSize	38
Description of com.odi.applicationName	39
Description of com.odi.cachedObjectCount.	39
Description of com.odi.cachedObjectTransactionAge	39
Description of com.odi.cacheSize	40
Description of com.odi.disableObjectCaching	40
Description of com.odi.disableWeakReferences.	41
Description of com.odi.migrateUnexportedStrings.	42
Description of com.odi.ObjectStoreLibrary	42
Description of com.odi.password and com.odi.user	42
Description of com.odi.product.	42
Description of com.odi.stringPoolSize	45
Description of com.odi.trapUnregisteredType	46
Description of com.odi.useImmediateStrings	46
Managing Databases	49
Creating a Database.	50
Method Signature for Creating a Database.	50
Example of Creating a Database.	50
Result of Creating a Database	51
Specifying a Database Name in Creation Method	52
When the Database Already Exists	52
Installing Schema on Database Creation	52
Creating Segments	53
Storing Objects in a Particular Segment	53
Iterating Through the Segments in a Database.	54

Creating Clusters	54
Storing Objects in a Particular Cluster	55
Iterating Through the Clusters in a Segment.	55
Determining Whether a Database, Segment, or Cluster Is Transient	
56	
Opening and Closing a Database.	56
Opening a Database.	57
Possible Open Modes	57
Opening the Same Database Multiple Times	58
Closing a Database	58
Automatic Opens of a Database	60
Objects in Closed Databases.	61
Moving or Copying a Database	61
Performing Garbage Collection in a Database	61
Background About the Persistent Garbage Collector.	61
API for Collecting Garbage in a Database	62
API for Collecting Garbage in a Segment	63
Command-Line Utility for Collecting Garbage	64
Running <code>osgc</code> on C++ Databases or Segments	64
Performing Compaction on a Database	64
Schema Evolution: Modifying Class Definitions of Objects in a	
Database	65
When Is Schema Evolution Required?	66
Preparing to Use the Schema Evolution API	67
Using the Schema Evolution API	67
Considerations for Using Serialization to Perform Schema Evolution	68
Steps for Using Sample Schema Evolution Serialization Code	69
Sample Code for Using Serialization to Perform Schema Evolution.	70
Dumping and Loading Databases	73
Destroying a Database	75
Obtaining Information About a Database	75
Is a Database Open?	75
What Kind of Access Is Allowed?	76
What Is the Pathname of a Database?	76

What Is the Size of a Database?	76
With Which Session Is the Database or Segment Associated?	77
Which Objects Are in the Database?	77
Are There Invalid References in the Database?	77
Which Databases Are Affiliated with This Database?	77
Grouping Objects in Multiple Clusters and Segments	77
Planning for Cross-Cluster and Cross-Segment References	78
Export Exceptions	80
How Many Exported Objects Are Needed?	80
Database Operations and Transactions	81
Upgrading Databases for Use with the JDK 1.2	83
Working with Transactions	85
Starting a Transaction	86
Calling the begin() Method	86
Allowing Objects to Be Modified in a Transaction	86
Difference Between Update and Read-Only Transactions	87
Working Inside a Transaction	87
Obtaining the Session Associated with the Current Transaction	87
Transaction Already in Progress	88
Obtaining Transaction Objects	88
Performing a Transaction Checkpoint	89
Setting a Transaction Priority	89
Ending a Transaction	89
Committing Transactions	90
What Can Cause a Transaction Commit to Fail?	91
Aborting Transactions	91
Handling Automatic Transaction Aborts	92
Results of Transaction Abort	93
Description of Transaction Abort Exceptions	93
Restarting Aborted Transactions	94
Handling Deadlocks	96
Determining Transaction Boundaries	96
Inconsistent Database State	96
Combining Transactions	97

Multiple Cooperating Threads	97
Performance Considerations	98
Storing, Retrieving, and Updating Objects	99
Storing Objects.	100
How Objects Become Persistent	100
Storing Objects in a Particular Segment or Cluster	101
What Is Reachability?.	101
Situations to Avoid	101
Storing Java-Supplied Objects.	102
Retrieving Persistent Objects	102
Steps for Retrieving Persistent Objects.	103
Determining the Database That Contains an Object.	103
Determining Whether an Object Has Been Stored	103
Locking Objects.	103
Working with Database Roots.	104
Creating Database Roots	104
Retrieving Root Objects	105
Roots with Null Values	106
Using Primitive Values as Roots.	106
Changing the Object Referred to by a Database Root.	106
Destroying a Database Root	107
Destroying the Object Referred to by a Database Root.	107
How Many Roots Are Needed in a Database?	107
Iterating Through the Objects in a Cluster, Segment, or Database	108
Using External References to Stored Objects	109
Creating External References	110
Obtaining Objects from External References	112
Encoding External References as Strings	112
Using the ExternalReference Field Accessor Methods	113
External References and Exported Objects	114
External Reference Equality	115
Reusing External Reference Objects.	115
External Reference Examples	116

Referencing Objects Stored in Other Databases.	117
Updating Objects in the Database	118
Background for Specifying Object State.	119
About Object Identity	119
About the Object Table	122
Committing Transactions to Save Modifications	122
Setting a Default Commit Retain State for a Session.	124
Setting Persistent Objects to a Default State	125
Making Persistent Objects Stale	125
Making Persistent Objects Hollow	126
Retaining Persistent Objects as Readable.	127
Retaining Persistent Objects as Writable	129
Retaining Persistent Objects as Transient	130
The Way Transient Fields Are Handled.	131
Caution About Retaining Unexported Objects	131
Evicting Objects to Save Modifications	132
Description of Eviction Operation	133
Setting the Evicted Object to Be Stale.	134
Setting the Evicted Object to Be Hollow.	134
Setting the Evicted Object to Be Read-Only	135
Summary of Eviction Results for Various Object States	135
Evicting All Persistent Objects	136
Evicting Objects When There Are Cooperating Threads	136
Committing Transactions After Evicting Objects	137
Evicting Objects Outside a Transaction	137
Aborting Transactions to Cancel Changes	138
Setting a Default Abort Retain State for a Session	139
Setting Persistent Objects to a Default State	140
Specifying a Particular State for Persistent Objects	140
Destroying Objects in the Database	142
Calling <code>ObjectStore.destroy()</code>	142
Destroying Objects That Refer to Other Objects	142
Destroying Objects That Are Referred to by Other Objects.	145
Default Effects of Various Methods on Object State	146

Transient Fields in Persistence-Capable Classes	146
Behavior of Transient Fields	147
Preventing fetch() and dirty() Calls on Transient Fields	147
Avoiding finalize() Methods	148
Troubleshooting Access to Persistent Objects	148
Handling Unregistered Types	149
How Can There Be Unregistered Types?	150
Can Applications Work When There Are Types Not Registered? . . .	150
What Does ObjectStore Do About Unregistered Types?	150
When Does ObjectStore Create UnregisteredType Objects?	151
Can Your Application Run with UnregisteredType Objects?	152
Troubleshooting ClassCastExceptions Caused by Unregistered Types . .	153
Troubleshooting the Most Common Problem	154
Troubleshooting OutOfMemoryError	154
Working with Collections	157
Description of ObjectStore Utility Collections	157
Introduction to java.util Interfaces and Classes	158
Description of OSHashBag	160
Description of OSHashMap	160
Description of OSHashSet	160
Description of OSHashtable	161
Description of OSTreeMapxxx	161
Description of OSTreeSet	163
Description of OSVector	165
Description of OSVectorList	165
Advantages of Using ObjectStore Utility Collections	166
Querying Collection Views of Map Entries	166
Background About Utility Collections and JDK 1.2 Collections	166
The Way to Choose a Collection	168
Comparing Collection Classes	170
Performance-Based Recommendations for Collections	170
Using ObjectStore Utility Collections	171
Creating Collections	171

Navigating Collections with Iterators	171
Performing Collection Updates During Iteration	172
Querying ObjectStore Utility Collections	173
Creating Queries.	173
Description of Query Syntax.	177
Sample Program That Uses Queries	179
Matching Patterns in Query Strings	179
Using Free Variables in Queries	181
Executing Queries.	182
Limitations on Queries.	183
Enhancing Query Performance with Indexes	184
How Indexes Work	184
Adding Indexes to Collections.	185
Dropping Indexes from Collections	186
Using Multistep Indexes in Queries	186
Sample Program That Uses Indexes	187
Sample Program That Queries User-Defined Fields	187
Modifying Index Values	187
Managing Indexes and Index Values	189
Optimizing Queries for Indexes	190
Manipulating Indexes Outside the Query Facility.	192
Storing Objects as Keys in Persistent Hash Tables	192
Requirements for Hash Code Methods.	192
Providing an Appropriate Persistent Hash Code Method	193
Storing Built-In Types as Keys in Persistent Hash Tables	193
Using Third-Party Collections Libraries	194
Generating Persistence-Capable Classes Automatically	195
Overview of the Class File Postprocessor	196
Description of the Annotations	197
Description of the Process	197
Postprocessing a Batch of Files Is Important	198
Postprocessor API	199
Manual Annotation	199

Running the Postprocessor	200
Preparing to Run the Postprocessor	200
Requirements for Running the Postprocessor	201
Example of Running the Postprocessor	202
About the Postprocessor Destination Directory	202
How the Postprocessor Interprets File Names	203
Order of Processing	204
How the Postprocessor Handles Duplicate File Specifications	205
How the Postprocessor Handles Files Not Found	206
.Zip and .Jar Files as Input to the Postprocessor	206
How the Postprocessor Handles Previously Annotated Classes	206
Troubleshooting OutOfMemory Error	206
How the Postprocessor Handles Inner Classes	207
When ClassInfo.java Files Are Generated	207
Managing Annotated Class Files	208
Ensuring That the Compiler Finds Unannotated Class Files	209
Ensuring That ObjectStore Finds Annotated Class Files	210
Using the Right Class Files in Complex Applications	210
Alternatives for Finding the Right Files	211
How the Postprocessor Determines Whether to Generate an Annotated Class File	211
Creating Persistence-Aware Classes	212
Specifying the Postprocessor Command Line	212
No Changes to Superclasses	213
How the Postprocessor Works	213
Ensuring Consistent Class Files	213
Modifications to Superclasses	214
Effects on Inheritance	214
Location of Annotated Class Files	215
Postprocessor Errors and Warnings	215
Handling of final Fields	215
Handling of Static Fields	216
Which Java Executable to Use	217
Line-Number and Local-Variable Information	217

Using a Debugger	217
Handling of finalize() Methods	218
Description of Postprocessor Optimizations	218
Including Transient and Already Annotated Classes	219
Copying Classes to the Destination Directory	220
Specifying Classes to Be Copied and Classes to Be Persistence Capable	220
When Can a Class Be Transient?	220
Putting Processed Classes in a New Package	221
Using the -translatepackage Option	221
How the Postprocessor Applies the Option	222
Updating References to New Package Name.	223
References to Transient and Persistent Versions of a Class	223
References to Transient Instances of a Persistence-Capable Class	224
Creating Persistence-Capable Classes with Transient Fields.	224
Transient Fields and Serialization	225
Initialization of Some Transient Fields	225
Customizing Updated Classes	226
Implementing Customized Methods and Hook Methods	226
Creating a Hollow Object Constructor	229
Optimizing Operations That Retrieve Persistent Objects	230
Procedure for Optimizing Operations	230
Inlining Code	231
Preventing Fetch of Transient Fields	231
Performing a Test Run of the Postprocessor	232
Using an Input File.	233
Annotations You Must Add	234
Interfacing with Nonpersistent Methods.	234
Interfacing with Native Classes.	234
Annotating Subclasses.	235
Passing Arrays	235
Implementing the Hollow Object Constructor for Some Instance Fields	235
Using the Java Reflection API with Persistence-Capable Objects	236

Class File Postprocessor Limitations	236
Generating Persistence-Capable Classes Manually.	237
Explicitly Defining Persistence-Capable Classes.	238
Implementing the IPersistent Interface	238
Defining the Required Fields	239
Defining Required Methods in the Class Definition	239
Implementing the IPersistentHooks Interface	241
Making Object Contents Accessible	242
Defining a ClassInfo Subclass	243
Example of a Manually Annotated Persistence-Capable Class	244
Additional Information About Manual Annotation	247
Defining a hashCode() Method	247
Defining a clone() Method.	248
Working with Transient-Only and Persistent-Only Fields	248
Defining Persistence-Aware Classes	251
Following Postprocessor Conventions	252
Annotating Abstract Classes	252
Creating and Accessing Fields in Annotations	253
Making Persistent Objects Accessible	253
Creating Fields	254
Getting and Setting Generic Object Field Values	255
Methods for Creating Fields and Accessing Them in Generic Objects	256
Controlling Concurrency.	259
Reducing Wait Time for Locks.	259
Clustering.	259
Transaction Length	260
Multiversion Concurrency Control (MVCC).	260
Lock Timeouts.	260
Conflicts Caused by Schema Installation.	261
Using Multiversion Concurrency Control (MVCC)	261
When Is MVCC Appropriate?	261
How Does MVCC Work?	261
Obtaining Read Locks.	262

Accessing Multiple Databases in a Transaction	262
Serializability	262
Opening a Database for MVCC Access	262
Determining Whether a Database Is Opened for MVCC	263
Updating the Snapshot	263
Where to Find Additional Information	264
Checkpoint: Committing and Continuing a Transaction	264
Advantages of a Checkpoint	265
Setting a Default Checkpoint Retain State for a Session	266
Setting Persistent Objects to a Default State	267
Locking Objects, Clusters, Segments, and Databases to Ensure Access	267
Description of Acquire Lock Methods	268
Locking Objects for Read or Write Access	269
Specifying the Wait Time for a Lock	269
Releasing Locks	270
Locking Peer Objects	270
Obtaining Information About Concurrency Conflicts	270
Setting the Client Name	271
Helping Determine the Transaction Victim in a Deadlock	271
Installing Schema Information in Batch Mode	271
Background About Schema Information	272
Procedure for Installing Schema in Batch	273
Identifying the Application Types	273
Creating a Database with Batch Schema Installation	275
Installing Application Types in the Database Schema	276
If You Do Not Run the Postprocessor	277
Using the Notification Facility	279
Background About How Notification Works	279
What Is a Notification?	280
What Is the Flow of a Notification?	280
Threads and Notifications	281
Transactions and Notifications	281
Security	282

Creating Notifications	282
Descriptions of Constructors	282
Retaining References to Persistent Objects	283
Maximum Data Lengths	284
Restriction on data Argument Content	284
Subscribing to Receive Notifications	284
Discarding Subscriptions	285
Unsubscribing from Notifications	285
Sending Notifications.	285
Retrieving Notifications	286
Reading Notifications.	287
Managing the Notification Process.	287
Notification Queue	287
Performance Considerations	288
Network Service	289
Using the Java Dynamic Data (JDD) Classes.	291
An Overview of JDD	291
Types.	292
Attributes	292
Entities.	293
Basic JDD Tasks	294
Defining Types and Their Attributes	294
Creating Entities of a Type	295
Querying a Type	296
A Simple JDD Application	297
Relationships	299
One-to-One Relationships.	300
One-to-Many Relationships.	300
Many-to-Many Relationships	301
Linked Objects and Many-to-Many Relationships	302
Relationship Example.	303
Improving Query Performance with Superindexes.	306
Queries and Default Indexing	306
Superindexing.	306

	Mixing Java Objects with JDD	307
	When You Must Use the ObjectStore API	308
	Pros and Cons of Mixing the ObjectStore API with JDD	308
	Using Extended JDD Classes	308
	Miscellaneous Information.	311
	Java-Supplied Persistence-Capable Classes.	311
	Description of Java-Supplied Persistence-Capable Classes	311
	Can Other Java-Supplied Classes Be Persistence Capable?.	313
	Description of Special Behavior of String Literals	315
	Example of String Behavior	316
	Destroying Strings	316
	Serializing Persistent Objects	317
	Using Persistence-Capable Classes in a Transient Manner.	319
	Comparing Java and C++ Persistent Storage Layouts	319
	Differences Between C++ and Java Interfaces to ObjectStore.	321
	Timing of the Write Lock Acquisition	321
	Opening the Same Database Multiple Times.	321
	Environment Variables	322
	Tools Reference	323
	osjcfp: Running the Postprocessor	323
	Postprocessor API	332
	osjcgcn: Generating Peer Classes	333
	Command-Line Format	333
	Additional Options.	334
	Example of Running the Peer Generator Tool	337
	osjcheckdb: Checking References in a Database	338
	osjshowdb: Displaying Information About a Database	339
	osjuphsh: Upgrading String Hash Codes in Databases	340
	osjversion: Obtaining ObjectStore Version Information	341
Appendix	Packaging Your Application for End Users	343
	Glossary	345
	Index	349

Preface

This chapter contains the following sections:

- Purpose
- Audience
- Scope
- The Way This Book Is Organized
- Notation Conventions
- ObjectStore on the World Wide Web
- Your Comments

Purpose

The *Java API User Guide* provides information and instructions for using the Java interface to ObjectStore (OSJI). The Java interface allows you to write Java applications that store and retrieve data in ObjectStore databases.

A companion book, the *Java API Reference*, is available on line. It provides detailed information about the classes provided with OSJI.

Audience

This book is for experienced Java programmers who want to write applications that use the Java interface to ObjectStore.

Scope

This book does not provide information for developing ObjectStore applications that use C++ and Java. See *Developing Java Applications That Access C++*.

This book supports Release 6.1 of the Java interface to ObjectStore. See the *Release Notes* for specific ObjectStore release numbers.

The Way This Book Is Organized

This book is organized as follows:

- Chapter 1, *Introducing ObjectStore*, on page 1, describes what ObjectStore does, shows the application architecture, and defines some important terms.
- Chapter 2, *Example of Using ObjectStore*, on page 13, describes the components your application must include to use ObjectStore.
- Chapter 3, *Using Sessions to Manage Threads*, on page 19, discusses how to initialize threads to use ObjectStore and how to use threads with ObjectStore sessions.
- Chapter 4, *Managing Databases*, on page 49, provides instructions for creating, opening, closing, garbage collecting, and upgrading databases.
- Chapter 5, *Working with Transactions*, on page 85, describes how to start and end transactions.
- Chapter 6, *Storing, Retrieving, and Updating Objects*, on page 99, discusses the steps for storing, retrieving, and updating data.
- Chapter 7, *Working with Collections*, on page 157, provides information about how to create collections of objects and run queries over the collections.
- Chapter 8, *Generating Persistence-Capable Classes Automatically*, on page 195, describes how to use the class file postprocessor to create persistence-capable classes.
- Chapter 9, *Generating Persistence-Capable Classes Manually*, on page 237, describes how to manually annotate classes you define so they are persistence capable.
- Chapter 10, *Controlling Concurrency*, on page 259, describes the APIs you can use to ensure your application's access to data.
- Chapter 11, *Using the Notification Facility*, on page 279, discusses the system you can set up to notify sessions when a previously defined event has taken place.
- Chapter 12, *Using the Java Dynamic Data (JDD) Classes*, on page 291, describes how to use the JDD classes to model and store dynamic data without having to run the postprocessor or perform schema evolution.

- Chapter 13, Miscellaneous Information, on page 311, discusses serialization, `String` literals, storage layout, differences from the C++ interface, and Java-supplied persistence-capable classes.
- Chapter 14, Tools Reference, on page 323, provides reference information for the ObjectStore utilities: `osgc`, `osjcfp`, `osjcgcn`, `osjcheckdb`, `osjshowdb`, and `osjversion`.
- Appendix, Packaging Your Application for End Users, on page 343, provides instructions for the files that you must include when you distribute your application.

Notation Conventions

This document uses the following conventions:

<i>Convention</i>	<i>Meaning</i>
Courier	Courier font indicates code, syntax, file names, API names, system output, and the like.
Bold Courier	Bold Courier font is used to emphasize particular code.
<i>Italic Courier</i>	<i>Italic Courier font</i> indicates the name of an argument or variable for which you must supply a value.
Sans serif	Sans serif typeface indicates the names of user interface elements such as dialog boxes, buttons, and fields.
<i>Italic serif</i>	In text, <i>italic serif typeface</i> indicates the first use of an important term.
[]	Brackets enclose optional arguments.
{ <i>a</i> <i>b</i> <i>c</i> }	Braces enclose two or more items. You can specify only one of the enclosed items. Vertical bars represent OR separators. For example, you can specify <i>a</i> or <i>b</i> or <i>c</i> .
...	Three consecutive periods indicate that you can repeat the immediately previous item. In examples, they also indicate omissions.

ObjectStore on the World Wide Web

ObjectStore has its own Web site (www.objectstore.net) that provides a variety of useful information about products, news and events, special programs, support, and training opportunities.

Technical Support

When you purchase technical support, the following services are available to you:

- You can send questions to support@objectstore.net. Remember to include your site ID in the body of the electronic mail message.
- You can call the Technical Support organization to get help resolving problems.
- You can access the Technical Support Web site, which includes
 - A template for submitting a support request. This helps you provide the necessary details, which speeds response time.
 - Frequently asked questions (FAQs) that you can browse and query.
 - Online documentation for all ObjectStore products.
 - White papers and short articles about using ObjectStore products.
 - Sample code and examples.
 - The latest versions of ObjectStore products, service packs, and publicly available patches that you can download.
 - Access to an ObjectStore product matrix.
 - Support policies.
 - Local phone numbers and hours when support personnel can be reached.

Education Services

Use the ObjectStore education services site (www.objectstore.net/services/education) to learn about the standard course offerings and custom workshops.

If you are in North America, you can call 1-800-477-6473 x4452 to register for classes. For information on current course offerings or pricing, send e-mail to classes@progress.com.

Searchable Documents

In addition to the online documentation that is included with your software distribution, the full set of product documentation is available on the ObjectStore Support Web server. The documentation is found at www.objectstore.net/documentation, and is listed by product. The site supports the most recent release and the previous supported release of ObjectStore documentation. Service Pack README files are also included to

provide historical context for specific issues. Be sure to check this site for new information or documentation clarifications posted between releases.

Your Comments

ObjectStore product development welcomes your comments about its documentation. Send any product feedback to support@objectstore.net. To expedite your documentation feedback, begin the subject with `Doc:`. For example:

Subject: Doc: Incorrect message on page 76 of reference manual

Chapter 1

Introducing ObjectStore

ObjectStore provides an application programming interface (API) that allows you to store Java objects persistently.

Contents	This chapter discusses the following topics:	
	What Is ObjectStore?	1
	What ObjectStore Does	2
	Benefits of Using ObjectStore	3
	Description of ObjectStore Process Architecture	3
	Definitions of ObjectStore Terms	5
	Prerequisites for Using the ObjectStore Java Interface	11

What Is ObjectStore?

ObjectStore is an object-oriented database management system. It allows you to

- Manipulate information in the database transparently by creating and modifying persistent Java objects
- Store and access data in the format in which it exists in the application
- Describe, store, and query complex data used in sophisticated software applications, as well as data traditionally managed by relational database applications
- Persistently store data independent of the data type

The Java interface to ObjectStore Development Client (referred to as ObjectStore) is for Java and C++ applications that require multiuser high-performance persistent storage for large databases with enterprise database features such as failover, on-line backup, fine-grained concurrency, and security.

ObjectStore can work well for distributed databases of virtually unlimited sizes and unlimited numbers of objects. ObjectStore supports the following:

- Applications that interface with databases and servers on local or remote machines
- Java applications and C++ applications that can access the same data
- Multiple concurrent sessions
- Multiple concurrent users
- Collections with indexed look-ups and queries
- Databases consisting of multiple segments that can have multiple clusters, which are variable-sized regions of disk space that ObjectStore uses to group objects stored in the database
- Operating on multiple databases in a transaction
- Cross-database references
- On-line backup, failover, and archive logging

ObjectStore PSE Pro for Java is a personal storage edition of the Java interface to ObjectStore. It is possible to use multiple ObjectStore Java products in the same Java virtual machine (VM). The `com.odi.product` property allows you to do this.

ObjectStore includes the `com.odi.odmg` package, which provides an Object Data Management Group (ODMG) binding. This binding includes classes for `Database` and `Transaction` that closely follow the ODMG specification. The package also includes the `com.odi.odmg.Collection` interface, persistence-capable classes that implement the `Collection` interface, and ODMG exception classes. See the `com.odi.odmg` package in the *Java API Reference*.

What ObjectStore Does

ObjectStore provides an API that allows a program to

- Start and end sessions to allow threads to use the ObjectStore API

- Create, open, close, and destroy databases
- Start, commit, and abort transactions to access data in the database
- Read and write database roots, which provide starting points for navigating to persistent objects
- Store objects in a database and retrieve and update those objects

ObjectStore can recover from an application failure or system crash. If a failure prevents some of the changes in a transaction from being saved to disk, ObjectStore ensures that none of that transaction's changes are saved in the database. When you restart the application, the database is consistent with the way it was before the transaction started.

With ObjectStore's archive logging facility, you can protect against media failure. See *Managing ObjectStore*.

Benefits of Using ObjectStore

ObjectStore provides a convenient and complete API for storing and sharing Java objects among users, hosts, and programs. After you define persistence-capable classes (classes whose instances can be stored in a database), writing an ObjectStore application is like writing any other Java application.

ObjectStore allows you to quickly read or modify portions of your persistent data. You are not required to read in all persistent data when you just want to look at a subset. This reduces start-up and transaction commit times and allows you to run much larger Java applications without increasing the amount of memory or swap space on the system.

When you access persistent data inside a transaction, ObjectStore ensures that your results are not compromised by other users sharing the data. If something goes wrong, or if you determine that you do not want to keep changes, you can abort the transaction. In that case, ObjectStore restores the database to the state it was in before the transaction started. This makes recovering from exceptions or failures straightforward.


Description of ObjectStore Process

Architecture

Following are the three kinds of processes in the ObjectStore environment:

- The *server* is the ObjectStore process that controls object storage. The server can manage databases for multiple client applications, which might be on multiple hosts.
- The *client* is the process in which ObjectStore links the ObjectStore client library into each ObjectStore application. In this way, each ObjectStore application is an ObjectStore client.
- The *cache manager* facilitates concurrent access to data by handling callback messages from the server to client applications.

The process architecture for the Java interface to ObjectStore is the same as for the C++ interface to ObjectStore. The server and the cache manager are the same regardless of the language interface you use with ObjectStore. See *Managing ObjectStore* for detailed information about managing these processes. The following figure shows the process architecture:



In the Java interface to ObjectStore, the client is a single process that has several software components. These components call each other as needed:

- ObjectStore C++ client library, which, among other things, communicates with the server

- ObjectStore Java client library, which provides the ObjectStore Java API
- Your Java application

The client library and the cache manager work together to maintain local copies of data on the client machine. Because local access is much faster than remote access, application performance is improved. ObjectStore ensures that you can never retrieve stale data, so you obtain performance benefits without sacrificing accuracy.

Definitions of ObjectStore Terms

This section describes the following terms, which you must be familiar with to use ObjectStore:

- Session
- Persistence Capable
- Persistent Object
 - Hollow persistent objects
 - Active persistent objects
 - Stale persistent objects
- Persistence Aware
- Primary Object
- Peer Object
- Transient Object
- Transitive Persistence
- Annotations
- Database Roots

Session

A *session* allows the use of the ObjectStore API. ObjectStore uses the abstract `com.odi.Session` class to represent sessions.

Your application must create a session before it can use any of the ObjectStore API. After a session is created, it is an active session. A session remains active until your application or ObjectStore terminates it. After a session is terminated, it is never used again. You can, however, create a new session.

A session creates a context in which you can create a transaction, access a database, and manipulate persistent objects. A session consists of a set of persistent objects and a set of ObjectStore API objects such as a `Transaction`, `Databases`, and `Segments`. In a single Java VM process, PSE Pro and ObjectStore allow multiple concurrent sessions.

Separate Java virtual machines can each run their own sessions at the same time. In addition, if you are using PSE Pro or ObjectStore, separate Java virtual machines can each run multiple sessions at the same time. See [How Sessions Keep Threads Organized](#) on page 19.

Persistence Capable

The term *persistence capable* refers to the capacity of an object to be stored in a database. If you can store the instances of a class in a database, the class is a persistence-capable class and the instances are persistence-capable objects.

The definition of a persistence-capable class includes specific annotations required by ObjectStore. After you compile class definitions, you run the ObjectStore class file postprocessor on the compiled classes to add the annotations that make the classes persistence capable. For more information, see [Chapter 8, Generating Persistence-Capable Classes Automatically](#), on page 195. In unusual circumstances, you might choose to add the annotations to the Java source file manually. For more information, see [Chapter 9, Generating Persistence-Capable Classes Manually](#), on page 237.

You must explicitly postprocess or manually annotate each class that you want to be persistence capable. The capacity for an object to be stored in a database is not inherited when you subclass a persistence-capable class.

Some Java-supplied classes are persistence capable. Other classes are not and cannot be made persistence capable. A third category of classes can be made persistence capable, but there are important issues to consider when you do so. Be sure to read [Java-Supplied Persistence-Capable Classes](#) on page 311.

Persistent Object

A *persistent object* is a representation of an object that is stored in a database. After an application retrieves an object from the database, the application works with the persistent object in the Java environment.

A persistent object always exists in one of three states:

- Hollow
- Active

- Stale

Hollow persistent objects

Methods you call can change the state of a persistent object.

A *hollow persistent object* has the same structure as the object in the database that it represents. A hollow object contains the same fields as the object in the database that the persistent object represents, but the fields of the hollow object usually do not contain values corresponding to those values stored in the database. The values for the fields in a hollow object might be `null` or the values from a previous transaction.

When your application acquires a reference to an object that has not yet been read in from the database, ObjectStore generates a hollow object as a placeholder for that object. ObjectStore does not actually read in the contents of the object until your application tries to read, write, or invoke a method on the object.

When your application accesses a hollow object, ObjectStore turns it into an active persistent object. ObjectStore retrieves the contents of the object from the database and stores them in the fields of the hollow object, which makes it an active persistent object. This process is referred to as initialization. In most applications, this happens automatically, because the postprocessor inserts the required calls.

Obtaining an object from a database root always results in a hollow object. If you get the same root three times, the object it identifies is still hollow. You must access the object to make it active.

After an application accesses the contents of persistent objects, the objects that the persistent object references are hollow objects unless their contents were accessed previously. For example, suppose that you have the following class:

```
class A {
    B b;
}
```

When you obtain a reference to an instance of `A`, ObjectStore creates a hollow `A` object to represent that instance. When you read or update the instance of `A`, ObjectStore turns it into an active object and creates a hollow `B` object to represent the referred-to instance of `B`. If you then read or update the instance of `B`, ObjectStore makes that hollow object into an active object and creates hollow objects for any objects referred to by `B`.

Active
persistent
objects

A persistent object must be active before an application can read or update it. An *active persistent object* starts as an exact copy of the object that it represents in the database. The contents of an active object are available to be read by the application and might be available to be modified. If an active object is updated by the application, it is no longer identical to the object in the database that it represents.

When a persistent object is active, ObjectStore internally flags it as either clean or dirty. An active object is marked initially as clean when its contents are read into memory. At this point, ObjectStore recognizes that the contents of the persistent object match the contents of the object in the database. An active object is dirty when it is a modified version of the stored object that the active object represents. When you modify an object, ObjectStore automatically changes the flag from clean to dirty. The class file postprocessor inserts the code that makes a persistent object clean or dirty.

For example, suppose that you have an instance of a `Person` object in which the `age` field has the value 30. When you read this object, it is in the clean state. If you modify the value of `age`, even if the new value you assign is 30, the object is then in the dirty state.

Stale
persistent
objects

A *stale persistent object* is no longer valid. Its fields have default values or values left over from previous transactions and should not be used. An object becomes stale when an application calls

- `Transaction.commit()` on the transaction in which the object could be read or modified, and the call
 - Does not specify a `retain` argument, or it specifies `ObjectStore.RETAIN_STALE`
 - And `Transaction.setDefaultCommitRetain()` has not been set or it specifies `ObjectStore.RETAIN_STALE`
 - And `Transaction.setDefaultRetain()` has not been set or it specifies `ObjectStore.RETAIN_STALE`

(There is not an API that sets a default `retain` value for the `evict()` or the `destroy()` methods. You can only set a default `retain` value for the `abort()`, `checkpoint()`, and `commit()` methods.)

- `Transaction.abort()` on the transaction in which the persistent object could be read or modified, and the call
 - Does not specify a `retain` argument or it specifies `ObjectStore.RETAIN_STALE`

- `And Transaction.setDefaultAbortRetain()` has not been set or it specifies `ObjectStore.RETAIN_STALE`
- `And Transaction.setDefaultRetain()` has not been set or it specifies `ObjectStore.RETAIN_STALE`
- `Transaction.checkpoint()` on the transaction in which the persistent object could be read or modified, and the call
 - Does not specify a retain argument or it specifies `ObjectStore.RETAIN_STALE`
 - `And Transaction.setDefaultCommitRetain()` has not been set or it specifies `ObjectStore.RETAIN_STALE`
 - `And Transaction.setDefaultRetain()` has not been set or it specifies `ObjectStore.RETAIN_STALE`
- `ObjectStore.evict()` on the object and the call does not specify a retain argument, or it specifies `ObjectStore.RETAIN_STALE`
- `ObjectStore.destroy()` on the object

An object can also become stale when the transaction in which it was accessed is aborted because of a deadlock or another action that caused `AbortException` to be signaled.

If an application tries to read or update a stale object, `ObjectStore` signals `ObjectException`. An application must not invoke any instance method on a stale object.

Persistence Aware

If the methods of a class can operate on fields of persistent objects, but instances of the class itself are not persistence capable, the class is *persistence aware*.

Typically, if you want a class to be persistence aware, you run the postprocessor on it to put in the required annotations. See Chapter 8, *Generating Persistence-Capable Classes Automatically*, on page 195. Occasionally, you might choose to annotate the class manually to make it persistence aware. See Chapter 9, *Generating Persistence-Capable Classes Manually*, on page 237.

When a method accesses fields in a persistent object, `ObjectStore` checks to ensure that the data has been read from the database. This checking is done by calls that the postprocessor inserts in your code. These are the annotations

mentioned in the previous paragraph. The annotations are calls to the `ObjectStore.fetch()` or `ObjectStore.dirty()` method.

Every persistence-capable and persistence-aware class must have these annotations. Persistence-capable classes also include many other annotations.

A class must be persistence aware only if it directly accesses the fields of a persistence-capable object. This includes access of elements of a persistent array. If your persistence-capable classes have only private fields and do not return arrays that might be persistent, other classes can call methods on the persistence-capable object without being persistence aware.

Primary Object

A *primary object* is a persistence-capable Java object. A persistence-capable object is an object that can be stored in an ObjectStore database. You can use primary objects as if they were ordinary Java objects. The database correctly records their identities, classes, and field values. Primary objects do not extend the `CPlusPlus` class.

Peer Object

Peer objects provide a way for Java applications to use C++ objects. A peer object acts as a proxy for a particular C++ object. It has no data fields, so it does not hold any state that is represented by the data members of the corresponding C++ object.

However, a peer object provides object identity, which allows you to invoke a method on the corresponding C++ object. You can call Java methods on a peer object to invoke all public methods of the original C++ object. You can think of a peer object as a handle to a C++ object.

Each peer object identifies exactly one C++ object. Multiple peer objects can represent the same C++ object. For example, each element of a C++ array of classes is represented by a peer object.

A peer object is an instance of a Java peer class. All peer objects extend the `CPlusPlus` class. Information about using peer objects is in *Developing Java Applications That Access C++*.

Transient Object

A *transient object* is an object that is not already in a database.

Transitive Persistence

When an application commits a transaction, ObjectStore stores in the database any transient objects that can be reached transitively from any persistent object. This is the process of *transitive persistence*. Transient objects that are referenced by persistent objects become persistent when the transaction commits. For this to work, the transient objects must be persistence capable.

Annotations

The class file postprocessor annotates classes you define so that they are persistence capable. This means that the postprocessor makes a copy of your class files, overwrites your original class files or places them in a directory that you specify, and adds byte code instructions (*annotations*) that are required for persistence. Complete information about annotations is in Chapter 8, *Generating Persistence-Capable Classes Automatically*, on page 195.

Occasionally, you might want to annotate your code manually. Information you need to do this is in Chapter 9, *Generating Persistence-Capable Classes Manually*, on page 237.

Database Roots

A *database root* provides a way to associate a name with an object in a database. Applications use database roots to locate one or more persistent objects for performing queries or navigating to other persistent objects. When you make an object the value of a persistent database root, doing so establishes the object as persistent and makes the objects it refers to available for transitive persistence.

At any given time, a database root is either associated with one database or it is null. You can change the database with which a root is associated. Information about database roots is in *Working with Database Roots* on page 104.

Prerequisites for Using the ObjectStore Java Interface

To use the ObjectStore Java interface, you must

- Be an experienced programmer familiar with the Java language.
- Have the supported platform as defined in the top-level `README` file.
- Have available a supported Java platform. The compiler must conform to JavaSoft specifications. The Java VM must be among those supported by ObjectStore. You cannot use a supported compiler with an unsupported VM. See Requirements for Using This Release in the *ObjectStoreJava Interface for ObjectStore Release Notes*.
- Have installed ObjectStore server and the C++ interface to ObjectStore Development Client.

You should also be familiar with the information in *Managing ObjectStore*, which describes the ObjectStore architecture and provides information about ObjectStore server parameters, ObjectStore client environment variables, and ObjectStore utilities that you can use.

- If you plan to use ObjectStore to operate on both Java and C++ objects, you must also be an experienced C++ programmer.

Chapter 2

Example of Using ObjectStore

This chapter provides a simple example of a complete ObjectStore program. The code for this example is in the `com\odi\demo\people` directory provided with ObjectStore.

Contents	This chapter discusses the following topics:	
	Overview of Required Components	13
	Sample Code	14
	Before You Run the Program	17
	Running the Program	18

Overview of Required Components

The sample program stores information about a few people, then retrieves some of the information from the database and displays it. The program shows the components you must include in your application so that it can use ObjectStore. These components are

- Create a session. The example calls the `Session.create()` method to start a nonglobal session. See page 24.
- Join a thread to a session. The example calls the `Session.join()` method to associate this thread with the session. See page 30.
- Create or open a database. The example creates the `person.odb` database and uses the `db` variable to refer to it. See page 50.

- Start and commit transactions as needed. The example uses one transaction to store the objects in the database. It then uses a second transaction to retrieve the stored objects. See page 86.
- Create a database root, which provides a starting point for accessing objects in the database. The example creates a root with the name "Tim" and associates it with the `tim` instance of the `Person` class. See page 104.
- Store objects referenced by a root in the database. The example stores `sophie` and `joseph` in the database when the transaction is committed. See page 100.
- Use a database root to retrieve objects from a database and do something with them. The example starts a new transaction, retrieves `tim`, and displays a line of information. See Retrieving Persistent Objects on page 102.
- End the session. The example calls the `Session.terminate()` method. This closes the open database and shuts down `ObjectStore`. See page 26.

When you write an `ObjectStore` program, you write it as though classes are persistence capable. However, a program cannot store objects persistently until you run the `ObjectStore`-provided class file postprocessor. The postprocessor generates annotated versions of the class files. The annotated version of the class definition is persistence capable. You run the postprocessor after you compile the program and before you run the program.

Sample Code

```
package com.odi.demo.people;

// Import the com.odi package, which contains the API:
import com.odi.*;

public
class Person {

    // Fields in the Person class:

    String name;
    int age;
    Person children[];

    // Main:

    public static void main(String argv[]) {
        try {
```

```

String dbName = argv[0];

// The following line starts a nonglobal session and
// joins this thread to the new session. This allows the
// thread to use ObjectStore.
Session.create(null, null).join();

Database db = createDatabase(dbName);
readDatabase(db);
db.close();
}

// The following shuts down ObjectStore.
finally {
    Session.getCurrent().terminate();
}
}

static Database createDatabase(String dbName) {
    // Attempt to open and destroy the database specified on the
    // command line. This ensures that the program creates a
    // new database each time the application is called.

    try {
        Database.open(dbName, ObjectStore.UPDATE).destroy();
    } catch (DatabaseNotFoundException e) {
    }

    // Call the Database.create() method to create a new
    // database.

    Database db = Database.create(dbName,
ObjectStore.ALL_READ | ObjectStore.ALL_WRITE);
    // Start an update transaction:

    Transaction tr = Transaction.begin(ObjectStore.UPDATE);

    // Create instances of Person:

    Person sophie = new Person("Sophie", 5, null);
    Person joseph = new Person("Joseph", 1, null);
    Person children[] = {sophie, joseph};
    Person tim = new Person("Tim", 35, children);

    // Create a database root and associate it with
    // tim, which is a persistence-capable object.
    // ObjectStore uses a database root as an entry
    // point into a database.

    db.createRoot("Tim", tim);

    // End the transaction. This stores the three person
    // objects, along with the String objects representing
    // their names, and the array of children, in the database.

    tr.commit();
}

```

```

        return db;
    }

    static void readDatabase(Database db) {
        // Start a read-only transaction:
        Transaction tr = Transaction.begin(ObjectStore.READONLY);

        // Use the "Tim" database root to access objects in the
        // database. Because tim references sophie and joseph,
        // obtaining the "Tim" database root allows the program
        // also to reach sophie and joseph.
        Person tim = (Person)db.getRoot("Tim");
        Person children[] = tim.getChildren();
        System.out.print("Tim is " + tim.getAge() + " and has " +
            children.length + " children named: ");
        for (int i=0; i < children.length; i++) {
            String name = children[i].getName();
            System.out.print(name + " ");
        }
        System.out.println("");
        // End the read-only transaction.
        // This form of the commit method ends the accessibility
        // of the persistent objects and makes the objects stale.
        tr.commit();
    }

    // Constructor:
    public Person(String name, int age, Person children[]) {
        this.name = name; this.age = age; this.children = children;
    }

    public String getName() {return name;}
    public void setName(String name) {this.name = name;}
    public int getAge() {return age;}
    public void setAge(int age) {this.age = age;}
    public Person[] getChildren() {return children;}
    public void setChildren(Person children[]) {
        this.children = children;
    }

    // This class is never used as a persistent hash key, so
    // include the following definition. If you do not, then
    // when you run the postprocessor it is unclear whether or
    // not you intend to use the class as a hash code.
    // Consequently, the postprocessor inserts a hashCode
    // function for you. The following definition avoids this.
    public int hashCode() {
        return super.hashCode();
    }
}

```


Before You Run the Program

Before you can run the sample program, you must

- Add an entry to your `CLASSPATH` environment variable
- Compile the source file
- Run the postprocessor on the `.class` file

Adding an Entry to CLASSPATH

In your `CLASSPATH` environment variable, you already have two entries related to ObjectStore:

- One entry for the `osji.jar` file to use ObjectStore
- One entry for the `tools.jar` file to use the class file postprocessor and other database tools

Ensure that these `.jar` files are explicitly in your class path. An entry for the directory that contains them is not sufficient.

Another entry is required for you to be able to build and run the program. This entry names the ObjectStore installation directory and allows ObjectStore to locate the annotated class files when you run the program.

For example, on Windows, if you place the ObjectStore distribution in the `c:\odi\osji` directory, you need the following entries:

```
c:\odi\osji\osji.jar;c:\odi\osji\tools.jar;c:\odi\osji
```

On UNIX, if you place the ObjectStore distribution in `/usr/local/osji`, you need

```
/usr/local/osji/osji.jar:/usr/local/osji/tools.jar:/usr/local/osji
```

Compiling the Program

To compile the program, change to the `com\odi\demo\people` directory and enter

```
javac *.java
```

As output, the `javac` compiler produces the byte code class file `Person.class`.

Running the Postprocessor

You must run the class file postprocessor to make the `Person` class persistence capable. The postprocessor generates new annotated class files. After you run the postprocessor, your program uses the annotated class files rather than the original class files.

Ensure that the `bin` directory that contains the `osjcfp` executable is in your path, as noted in the `README` file in the installation directory and the postprocessor documentation. See *Preparing to Run the Postprocessor* on page 200.

On Windows, to run the postprocessor, enter

```
osjcfp -dest . -inplace Person.class
```

On UNIX, to run the postprocessor, enter

```
osjcfp -dest . -inplace Person.class
```

The `-dest` option specifies a destination directory for the annotated files. It is a required option. The `-inplace` option specifies that the postprocessor should overwrite the original class files. When you specify the `-inplace` option, the postprocessor ignores the `-dest` option.

The result from the `osjcfp` command shown above is the annotated class file `com\odi\demo\people\Person.class`.

The `-inplace` option is the best choice for this example. However, when you are in an iterative development cycle, it is best not to specify `-inplace`. During development, putting the postprocessed files in a different directory avoids errors.

Running the Program

Run the program as a Java application. Following is a typical command line:

```
java com.odi.demo.people.Person person.odb
```

The argument is the pathname of the database.

The expected output is

```
Tim is 35 and has 2 children named: Sophie Joseph
```

Also, the example application creates or replaces the `person.odb` database in the current directory.

Chapter 3

Using Sessions to Manage Threads

This chapter provides information about how to manage the threads in your application. Sample code that uses threads is in `com\odi\demo\threads`.

Contents

This chapter discusses the following topics:

How Sessions Keep Threads Organized	19
Creating Sessions	23
Working with Sessions	25
Associating Threads with Sessions	27
Working with Threads	31
Threads and Persistent Objects	34
Description of Concurrency Rules	36
Description of ObjectStore Properties	37

How Sessions Keep Threads Organized

For a thread to use ObjectStore, it must be associated with a session. To use threads with ObjectStore, you must create at least one session and you must understand how to work with sessions.

If you try to use the ObjectStore API and you have not created a session, ObjectStore signals `NoSessionException`.

This section discusses the following:

- What Is a Session?
- How Are Threads Related to Sessions?
- What Is the Benefit of a Session?
- What Kinds of Sessions Are There?

What Is a Session?

A session allows the use of the ObjectStore API. ObjectStore uses the abstract `com.odi.Session` class to represent sessions.

Your application must create a session before it can use any of the ObjectStore API. After a session is created, it is an active session. A session remains active until your application or ObjectStore terminates it. After a session is terminated, it is never used again. You can, however, create a new session.

A session creates a context in which you can create a transaction, access one or more databases, and manipulate persistent objects.

Concurrent sessions

In a single Java VM, PSE Pro and ObjectStore allow multiple concurrent sessions.

If you are using ObjectStore or PSE Pro, separate Java virtual machines can each run multiple sessions at the same time.

See Description of Concurrency Rules on page 36.

How Are Threads Related to Sessions?

At any given time, an active session has zero or more associated *threads*. Any number of threads can join a session. Each thread can belong to only one session at a time.

At any given time, each thread is either joined to a single session or not joined (not associated) to a session. A thread that is not associated with a session can join a session. A thread that is associated with a session can leave the session to end its association with that session. It can rejoin the session at a later time or it can join another session.

For a thread to use the ObjectStore API, it must automatically or explicitly be associated with a session. All threads that join the same session cooperate with each other. ObjectStore does not prevent cooperating threads from accessing the same object. Consequently, it is your responsibility to identify code segments that must be synchronized. To successfully call the

`Session.join()` method to join a session, a thread must not already be associated with any session.

The *current thread* is the thread that you are making a call from. The *current session* is the session that the current thread belongs to.

What Is the Benefit of a Session?

The benefit of a session is apparent when you want to have more than one session. Two sessions in the same Java process allow you to perform two distinct activities that involve ObjectStore. Each session has a clean, isolated view of the database. If you want to have two or more independent transactions going on at the same time, you can use two or more sessions. Concurrent sessions can be accessing the same database or different databases.

When two sessions are accessing the same object in the database, there are two distinct persistent objects. Each session has its own persistent object, which is a copy of the object in the database. At least initially, these two persistent objects have the same content.

Each session has its own set of persistent objects and API objects. In most circumstances, the threads of session A are not allowed to operate on the persistent objects of session B. An exception to this rule is described in Multiple Representations of the Same Object on page 34.

Independent threads

A need for many different independent transactions normally arises because you have many Java threads with different things occurring in each one. Typically, this happens with a multithreaded application server, in which there are many threads. Each thread serves a different client, so you might want to have many threads. Each thread runs a separate transaction, and each thread is separate from each other thread.

Cooperating threads

On the other hand, there are times when you have multiple threads that are cooperating on some database task and must operate on the same objects at the same time. In this case, you might want two different Java threads to participate in the same transaction.

Controlling the threads that cooperate

Sessions allow you to control the threads that cooperate in a transaction and work in independent transactions. A session groups together a set of cooperating threads. Each session has a sequence (in time) of transactions and a set of associated threads that participate in these transactions.

There is a many-to-one relationship between threads and sessions. That is, any number of threads can belong to one session.

Example of cooperating threads

A common case of cooperating threads arises when you are writing a Java applet. In an applet, there are calls to different parts of your program in different threads. You have to specify for ObjectStore that all these threads are part of the same session. This allows them to operate on the same objects and in the same transactions. A similar situation exists when you use RMI and CORBA servers; that is, there is a control mechanism that calls your methods in different threads.

What Kinds of Sessions Are There?

An active session can be a global session or a nonglobal session. ObjectStore provides two kinds of sessions because when you need only one session, ObjectStore can do many things for you automatically.

Joining threads to sessions

As mentioned earlier, before you can use ObjectStore, you must create a session. For a thread to use ObjectStore, it must join a session. In a global session, an unassociated thread that makes a call to the ObjectStore API joins the session automatically. In a nonglobal session, this happens only when the call implies the session. See *Rules for Joining a Thread to a Session Automatically* on page 29. Otherwise, you must explicitly add the thread to a session.

Number of sessions

An active global session is the only session in the Java VM. With PSE Pro and ObjectStore, you can have multiple nonglobal sessions or one global session in a Java VM.

If you are using ObjectStore and you want to have multiple sessions, you must call the `ObjectStore.setNumExpectedSessions` method before you create any sessions. Otherwise by default, you can create only one session in the Java VM. This method must be called before the first session is created or an `ObjectStoreException` is signaled.

If you do not specify a value for the `com.odi.addressSpaceSize` property, then the value you specify for the `ObjectStore.setNumExpectedSessions` method along with the values set for the `OS_DEFAULT_AS_PARTITION_SIZE` and `OS_AS_SIZE` environmental variables are used to determine the default amount of C++ client address space and C++ client cache space assigned to newly created sessions.

Global sessions

Global sessions make programming easier because you need not know the ObjectStore APIs for associating threads with sessions. All threads that make ObjectStore API calls join the one global session automatically.

The drawback is that you can have only one session. If you change your program in the future to use multiple sessions, you might have to go back

and put in API calls to associate threads with the appropriate session. If you think you might use multiple sessions in the future, it would probably be a good idea to prepare for that by using a nonglobal session and explicitly joining the session in each thread.

Creating Sessions

When you create a session, you initialize `ObjectStore` for use by the threads that become associated with that session. You can create a session in the following ways:

- Call the `Session.createGlobal()` method to create a global session.
- Call the `Session.create()` method to create a nonglobal session.

Regardless of how you create a session, you can specify a number of `ObjectStore` properties when you create the session. These properties determine how `ObjectStore` behaves in a variety of situations.

Creating Global Sessions

When the session is a global one and a thread that is not associated with the session calls an `ObjectStore` API, `ObjectStore` automatically joins the thread to the session. After you create a global session, you need not be concerned about joining threads to the session.

To create a global session, call the `Session.createGlobal()` method. The method signature is

```
public static Session createGlobal(String host,
    java.util.Properties properties)
```

This method creates and returns a new session and designates the session as a global session. No threads are joined to this session yet. Any thread, including the thread that creates the session, automatically joins the session the first time the thread uses `ObjectStore`.

`ObjectStore` ignores the first parameter; you can specify `null`. The second parameter specifies `null` or a property list. See [Description of ObjectStore Properties](#) on page 37.

If you try to create a global session when there is already an active session, `ObjectStore` signals `ObjectStoreException`.

To obtain a global session, call `Session.getGlobal()`. The method signature is

```
public static Session getGlobal()
```

If the global session is active, `ObjectStore` returns it. Otherwise, `ObjectStore` returns null.

Creating Nonglobal Sessions

You can create a nonglobal session. The difference between a global and nonglobal session is that in a nonglobal session

- `ObjectStore` does not join all unassociated threads to the session automatically.
- Multiple nonglobal sessions can exist in the same Java VM for `ObjectStore` and PSE Pro.

Joining threads to sessions

For `ObjectStore` to join an unassociated thread to a nonglobal session automatically, the thread must be making an `ObjectStore` API call that implies a session. See [Rules for Joining a Thread to a Session Automatically](#) on page 29.

You must explicitly join a thread to a session before that thread can call an `ObjectStore` API that does not imply a session. See [Explicitly Associating Threads with a Session](#) on page 30.

Method signature

The method signature for creating a nonglobal session is

```
public static Session create(String host,  
                             java.util.Properties properties)
```

This method creates and returns a new session. `ObjectStore` ignores the `host` argument; specify null. The second argument specifies null or a property list. See [Description of ObjectStore Properties](#) on page 37.

`ObjectStore` does not join the calling thread to the session. If you are using `ObjectStore` or PSE Pro, a thread can belong to a nonglobal session and call a method that creates another nonglobal session.

Session name

`ObjectStore` generates a name for the session and never reuses that name for the lifetime of the process in which the session was created. If you want to specify a particular name, use the following overloading to specify a unique session name:

```
public static Session create(String host,  
                             java.util.Properties properties,  
                             String name)
```


ObjectStore uses the session name in debugging messages. The `Session.getName()` method returns the name of the session.

Exception conditions

If you call `Session.create()` when there is an active global session, ObjectStore signals `ObjectStoreException`.

Working with Sessions

After you create a session, you need to know how the session functions with regard to transactions. You also need to know about the operations you can perform on the session. This section discusses

- Sessions and Transactions
- Shutting Down Sessions
- Obtaining a Session
- Determining Whether a Session Is Active

Sessions and Transactions

At any given time, a session has one associated transaction in progress or it does not have any associated transaction. Each transaction is associated with exactly one active session.

When a session is created, there is no associated transaction. While a session is active, an application can start, then commit or abort, one transaction at a time per session. Over time, a session is associated with a sequence of transactions.

If a transaction is in progress when an application or ObjectStore shuts down the session, ObjectStore aborts the transaction as part of the shutdown process.

Within a session, multiple databases can be open at the same time.

All transactions can access the same database.

See Description of Concurrency Rules on page 36.

Transaction in progress?

To determine whether there is a transaction in progress, call the `Session.inTransaction()` method. The method signature is

```
public boolean inTransaction()
```

If there is a transaction associated with the session, this method returns `true`. If there is no transaction associated with the session, this method returns `false`. If the session has been terminated, `ObjectStore` signals `NoSessionException`.

Obtaining the associated transaction

To obtain the transaction associated with a session, call the `Session.currentTransaction()` method. The method signature is

```
public Transaction currentTransaction()
```

If the session has been terminated, `ObjectStore` signals `NoSessionException`. If no transaction is associated with the session, `ObjectStore` signals `NoTransactionInProgressException`.

Obtaining the transaction's session

To obtain the session associated with a transaction, call the `Transaction.getSession()` method. The method signature is

```
public Session getSession()
```

Shutting Down Sessions

An application can shut down sessions. One reason you might want to shut down a session is to release the Java objects associated with the session. However, shutting down a session does not release all resources used by the C++ client. Such resources are released only when the application exits. To shut down a session, call the `Session.terminate()` method. The method signature is

```
public void terminate()
```

It does not matter whether the session is a global session or a nonglobal session; `ObjectStore` shuts down the session. If there are no other sessions, no thread can use the `ObjectStore` API until there is a new active session. The terminated session is never reused.

Transaction in progress

If the session you shut down has an associated transaction, `ObjectStore` aborts the transaction. If the session has already been terminated, `ObjectStore` does nothing. If the session has any associated threads, `ObjectStore` causes them to leave the session. If the session has an open database, `ObjectStore` closes it.

If `ObjectStore` signals `FatalException`, this shuts down the session.

Obtaining a Session

You can obtain a session with a call to any of the methods in the following list:

- `Placement.getSession()`
- `Session.getCurrent()`
- `Session.getGlobal()`
- `Session.of(object)`
- `Session.ofThread(thread)`
- `Transaction.getSession(thread)`

Determining Whether a Session Is Active

To determine whether a session is active, call the `Session.isActive()` method. The method signature is

```
public boolean isActive()
```

If the session is active, this method returns `true`. If the session has been terminated, this method returns `false`.

Associating Threads with Sessions

To help you associate threads with sessions, this section discusses

- Joining Threads to a Session Automatically
- Associating a Persistent Object with a Session
- Rules for Joining a Thread to a Session Automatically
- Examples of Calls That Imply Sessions
- Examples of Calls That Do Not Imply Sessions
- Explicitly Associating Threads with a Session

There is a bug in the software that prevents threads from being joined to sessions automatically. As a work around, you must explicitly join each thread to a session. See the release notes for details. This bug will be fixed in a future release.

Joining Threads to a Session Automatically

Whether a thread can join a session automatically depends on

- Whether the session is global or nonglobal
- Whether the API call that the thread is making implies a session

Global sessions

When there is a global session, an unassociated thread that makes a call to the ObjectStore API joins the global session automatically, if necessary. In the following situations, it might not be necessary to join the thread to the session:

- An unassociated thread calls a method on a transient object and the method requires a persistent object. Because the object is not persistent, the method cannot do anything, so it need not be joined to the session.
- An unassociated thread calls a method that does not operate on persistent objects, for example, calls to `ObjectStore.getAutoOpenMode()` and `ObjectStore.setLazyWriteLocking`.
- An unassociated thread calls a method that has already been executed. The thread might join the session automatically if it executes the method anyway. For example, when an unassociated thread tries to open a database that is already open, ObjectStore joins the thread to the session that the database belongs to, even though the thread does not actually do anything.

Nonglobal sessions

In a nonglobal session, ObjectStore automatically joins threads to the session when the call from the thread implies that session. This means that the call specifies an argument that is already associated with that session. This includes the object on which the method is invoked.

After ObjectStore automatically joins a thread to a session,

- The thread is associated with the session until you remove it from the session or the session terminates.
- ObjectStore performs the called method.

Associating a Persistent Object with a Session

How does an object become associated with a session? It happens implicitly. Assume that a thread is already associated with a session. This associated thread successfully calls an ObjectStore API. If there are any objects that result from that call, ObjectStore associates them with the session that the calling thread belongs to.

As a result of explicit and implicit association, a session provides a context for a set of persistent objects and a set of ObjectStore API objects, such as Database objects and a Transaction object.

The session defines a namespace. The namespace defines unique names (and, consequently, identities) for databases, segments, clusters, transactions, and persistent objects. While it is possible for threads in

different sessions to share objects, doing so is incorrect and usually results in exceptions.

If the thread in which an object was materialized leaves the session, the object remains associated with the session.

Rules for Joining a Thread to a Session Automatically

Because of the associations between objects and a particular session, some API calls imply a session. If there is no global session,

- A call that implies a session allows the calling thread to be joined to the implied session automatically.
- A call that does not imply a session does not allow the calling thread to be joined to a nonglobal session automatically.

If a thread associated with one session makes a call that implies another session, `ObjectStore` signals `WrongSessionException`.

Examples of Calls That Imply Sessions

A call that implies a session is a call that specifies an argument that is already associated with a session. It can also be a call in which the object on which the method is called is associated with a session. When these calls are in a thread that is not associated with a session, `ObjectStore` automatically joins the thread to the session with which the argument is already associated. It does not matter whether it is a global or nonglobal session. Some examples of API calls that imply a session follow:

- `Database.close()` — The `Database` argument was associated with a session when it was created.
- `ObjectStore.migrate(object, placement, export)` — The `placement` argument specifies a segment or database, which was associated with a session when it was initialized.
- `ObjectStore.destroy(object)` — The `object` argument designates a persistent object. It was associated with a session the first time it was accessed.

Examples of Calls That Do Not Imply Sessions

A call that does not imply a session is a call that does not specify an argument that is associated with a session. When these calls are in a thread that is not associated with a session, `ObjectStore` cannot join the thread to a session

automatically if it is a nonglobal session. Examples of API calls that do not imply a session follow:

- `Database.open(name, openType)` — A static method. The `Database` object does not exist yet so the `name` argument is not associated with a session.
- `Transaction.begin(type)` — Another static method, and the `Transaction` object does not exist yet.
- `ObjectStore.majorRelease()` — Also a static method.
- A call that accesses a transient object.
- A call that never accesses a persistent object, for example, `ObjectStore.getAutoOpenType()` and `ObjectStore.setLazyWriteLocking()`.

Explicitly Associating Threads with a Session

To join a thread to a session explicitly, you call the `Session.join()` method.

Session.join() To associate a thread with a session explicitly, call the `Session.join()` method. The method signature is

```
public void join()
```

This associates the current thread (the thread that contains the call to `join()`) with the session on which the `join()` method is called.

`ObjectStore` signals exceptions when joining threads to sessions in the following cases:

- When you try to join a thread to a session that has been terminated, `ObjectStore` signals `NoSessionException`.
- When you try to join a thread that already is in a session to another session, `ObjectStore` signals `WrongSessionException`.

However, if you try to join a thread to a session to which the thread already belongs, `ObjectStore` does nothing.

To join a thread to a session for a bounded duration of time, try something like the following:

```
Session session;  
session.create(null, myproperties);  
try {  
    session.join();  
    ...;  
    ...;
```

```

    ...;
} finally {
    session.leave();
}

```

Working with Threads

After you associate a thread with a session, it is important that you understand how to use the thread within the framework of a session. To that end, this section discusses

- Cooperating Threads
- Noncooperating Threads
- Synchronizing Threads
- Removing Threads from Sessions
- Threads That Create a Session
- Other Threads
- Determining Whether ObjectStore Is Initialized for the Current Thread

Cooperating Threads

All threads associated with a particular session cooperate with each other. That is, they

- Share transactions, persistent objects, and locks on ObjectStore data
- View the same state of any databases they access

For example, suppose thread *A* and thread *B* are cooperating threads (that is, they belong to the same session). *A* and *B* are running asynchronously. Each thread is issuing a sequence of operations and these sequences are interleaved in an unpredictable manner.

For ObjectStore, it is as if these operations are all coming from the same thread. It does not matter which operation comes from *A* and which operation comes from *B*. ObjectStore views the operations as being in a single sequence because they are issued from cooperating threads.

If *A* or *B* starts a transaction, it does not matter which thread issues the call. The transaction begins for both threads, regardless of the thread that actually starts the transaction. Any changes performed by *A* or *B* during the transaction are visible to both threads and can be acted on by either thread.

Similarly, if A commits the transaction, it is just as if B commits the transaction. So B must be in a state in which it is all right to commit the transaction. A and B must cooperate.

Noncooperating Threads

Threads that do not belong to the same session cannot share transactions, persistent objects, or locks on data, and cannot view the same state of the database. Threads that belong to different sessions are noncooperating threads. With ObjectStore or PSE Pro, a different session can belong to the same process or a different process.

Two or more noncooperating threads can open the same database at the same time and access the same root object. If two or more noncooperating threads access the same object in the database, an equivalent number of distinct instances of the persistent object exist — one for each thread. The identity test, `==`, does not show them to be identical.

Noncooperating threads can experience deadlock. See page 96.

Synchronizing Threads

Your application is responsible for synchronizing activity among cooperating threads when the transaction is committed or aborted. In general, your application must avoid accessing the database while a thread is committing the transaction and until a cooperating thread starts a new transaction. If a transaction is aborted, cooperating threads might need to retry database operations.

Additional information about synchronizing threads is in Multiple Cooperating Threads on page 97.

Removing Threads from Sessions

A thread can leave a session at any time, including while a transaction is in progress. This does not affect the transaction, nor any threads that are still joined to that session. With or without a transaction in progress, it is all right if no threads are associated with a session. The session does not terminate. A thread can join a session later to finish the transaction. If no thread does that, ObjectStore aborts the transaction when the session terminates.

To end the association of a thread with a session, call the `Session.leave()` method. The method signature is

```
public static void leave()
```


After you execute this method, the current thread is no longer joined to the session. If the current thread is already not associated with the session on which the method is called, `ObjectStore` signals `NoSessionException`.

If your application or `ObjectStore` shuts down a session, `ObjectStore` causes any associated threads to leave the session before it performs the shutdown.

If a thread is associated with a session and the thread terminates, it leaves the session automatically.

Threads That Create a Session

There is nothing special about the thread that creates a session. This thread can leave the session and any threads associated with that session can continue operating.

When your application calls the `Session.createGlobal()` or `Session.create()` method, `ObjectStore` does not associate the thread that calls the method with the newly created session. For that thread to join the new session, it must call the `Session.join()` method.

Other Threads

A thread that does not belong to a session cannot use the `ObjectStore` API. However, a thread need not be associated with a session to call `Session.isActive()` successfully.

If a session has a transaction in progress, a thread that is not associated with that session must not use persistent objects that belong to that session. See [Threads and Persistent Objects](#) on page 34.

If a session does not have a transaction in progress, any thread, including threads that do not belong to that session, can access persistent objects to the degree they were left visible when the application committed or aborted the transaction. See [Ending a Transaction](#) on page 89.

Determining Whether ObjectStore Is Initialized for the Current Thread

You can use the `Session.getCurrent()` method to determine whether `ObjectStore` is initialized for the current thread. The method signature is

```
public static Session getCurrent()
```

This method returns the session with which the current thread is associated. If the current thread is not associated with a session, this method returns `null`.

Threads and Persistent Objects

Each persistent object is associated with exactly one session. Any modification to the state of a persistent object must be done by a thread that cooperates in the session to which the persistent object belongs.

After you terminate a session, the persistent objects and API objects that were associated with it when it was terminated continue to be associated with the terminated session. One exception to this is when you call the `Database.close()` method with a `true` argument. This causes the persistent objects to be retained as transient objects, which are not associated with any session.

The information in this section is provided to help you ensure that threads access the correct objects. This section discusses these topics:

- Multiple Representations of the Same Object
- Example of Multiple Sessions
- Application Responsibility
- Effects of Committing a Transaction
- API Objects and Sessions

Multiple Representations of the Same Object

When you have multiple sessions, it is possible to have multiple persistent objects that represent the same object in the database. For example, a thread belonging to session A accesses object X. Then a thread belonging to session B accesses object X. There are two persistent objects that represent X. Each is a representation of the same object in the database. If you use the `==` operator on session A's X and session B's X, the result is that they are not identical; they are not the same object. Within a session, `ObjectStore` preserves object identity.

Example of Multiple Sessions

The following example shows actions you can and cannot perform when you have multiple sessions. In this example, suppose you have `sessionA` and `sessionB` and

- `threadA` is associated with `sessionA`.
- `threadB` is associated with `sessionB`.

In `threadA`, you start a transaction and read the contents of a persistent object called `objectA`. Because `threadA` is associated with `sessionA`, `objectA` belongs to `sessionA`. You commit the transaction with `ObjectStore.RETAIN_UPDATE`.

At this point, in `threadB`, you can read or modify `objectA` unless a transaction is in progress in `sessionB`. However, any modifications are discarded when `sessionA` starts a transaction.

Application Responsibility

It is the responsibility of the application to ensure that noncooperating threads act on persistent objects only in the ways allowed when a transaction is not in progress.

If you have a Java static variable that contains a persistent object and there are two separate sessions, you must decide the session that owns the static variable. In other words, if there is a Java static variable whose value is a persistent object, that persistent object is associated with one session.

Effects of Committing a Transaction

When a thread commits a transaction, it affects only those persistent objects that belong to the same session that the thread belongs to.

- | | |
|----------------------|---|
| Caution | You must ensure that an object never refers to an object that belongs to a different session. This is crucial because transitive persistence (performed when committing a transaction) must never reach an object that belongs to another session. If it does, <code>ObjectStore</code> signals <code>WrongSessionException</code> . |
| Array objects | When a thread commits a transaction, if <code>ObjectStore</code> reaches an object whose class does not implement <code>IPersistent</code> , <code>ObjectStore</code> treats the object as a transient object and migrates it to a database. This works correctly for immutable classes such as <code>Integer</code> and <code>String</code> . For array objects, this can cause unpredictable results because one session might modify the object while another session is using the old contents. |

API Objects and Sessions

Each ObjectStore API object is related to one session. These metaobjects are

- `Cluster`
- `Database`
- `DatabaseRootEnumeration`
- `DatabaseSegmentEnumeration`
- `Segment`
- `Transaction`

If you open the same database from two noncooperating transactions, each session has its own `Database` object to represent the database. These `Database` objects are not identical; that is, `==` returns false.

If you try to use a database in the wrong session, ObjectStore signals `WrongSessionException`.

Description of Concurrency Rules

ObjectStore allows multiple readers or one writer of an object at any given time. The term one writer implies one session in any process. In a session, two cooperating threads that are both updating an object count as one writer. Two threads from different processes do not cooperate and, therefore, count as two writers.

If you use Multiversion Concurrency Control (MVCC), there can be one writer and multiple readers. See Chapter 10, *Controlling Concurrency*, on page 259.

See *Handling Deadlocks* on page 96.

Granularity of Concurrency

ObjectStore locks data at the page level. ObjectStore acquires a read lock on the object the first time the Java application reads or writes the contents of the object in a transaction. When that happens, the underlying C++ ObjectStore acquires a read lock on all pages used to store the object, as well as any additional locks that are required on the schema information that is stored in the separate schema segment. There is also Java-specific metadata that is locked.

The size of a page is determined by your operating system.

Converting Read Locks to Write Locks

When a Java application modifies an object for which it already has a read lock, `ObjectStore` does not necessarily convert the read lock to a write lock immediately. The `ObjectStore.setLazyWriteLocking()` method controls this behavior. If lazy write locking is true (the default), `ObjectStore` acquires write locks only when it attempts to write the modified contents of the object to the database. That occurs either at commit time or when and if the application calls `ObjectStore.evict()` on the modified object.

Description of ObjectStore Properties

When you create a session, you can specify a properties argument. This section provides the following information about this argument:

- About Property Lists Relevant to `ObjectStore`
- Description of `com.odi.addressSpaceSize`
- Description of `com.odi.applicationName`
- Description of `com.odi.cachedObjectCount`
- Description of `com.odi.cachedObjectTransactionAge`
- Description of `com.odi.cacheSize`
- Description of `com.odi.disableObjectCaching`
- Description of `com.odi.migrateUnexportedStrings`
- Description of `com.odi.ObjectStoreLibrary`
- Description of `com.odi.password` and `com.odi.user`
- Description of `com.odi.product`
- Description of `com.odi.stringPoolSize`
- Description of `com.odi.trapUnregisteredType`
- Description of `com.odi.useImmediateStrings`

About Property Lists Relevant to ObjectStore

When you create a session, there are the following two relevant property lists:

- The `java.util.Properties` object that is the second argument to the method that creates the session
- The system property list

Finding the value of a property

To find the value of a property, ObjectStore checks the `java.util.Properties` object. If it provides a value, ObjectStore uses it. If it does not provide a value, ObjectStore checks the system property list.

Passing a property value

When you want to pass a property value to the method that creates a session, you typically put it in the `java.util.Properties` object that is an argument to the method that creates the session.

There is only one system property list for each Java VM. Multiple sessions in the same Java VM all use the same system property list. For more information about system properties, see `System.getProperty`, in section 21.6 of the *Java Language Specification*.

Defining a system property

All ObjectStore property names start with `com.odi`. You can pass in property information by defining it as a system property. For example:

```
Properties props = System.getProperties();
props.put("com.odi.useDatabaseLocking", "true");
Session session = Session.create(null,props);
```

There is also a `System.setProperties()` method that resets the System property list.

The JDK allows you to specify a system property by including

`-Dparameter=value`

on the `java` command line before the class name. Each such specification defines one system property. Not all Java virtual machines run this way.

Defining a Properties object

If you want to construct your own property list, the type of the property list argument is `java.util.Properties`. For example:

```
Properties props = new Properties();
props.put("com.odi.useDatabaseLocking", "true");
Session session = Session.create(null,props);
```

Description of `com.odi.addressSpaceSize`

The `com.odi.addressSpaceSize` property specifies in bytes the amount of C++ client address space available for the current session. The actual value used is the amount specified rounded up to the nearest 64 KB.

If the amount of address space requested is not available, an exception is signaled. If the `addressSpaceSize` property is not specified, the default amount of address space is determined by calls to `ObjectStore.setNumExpectedSessions()`, as well as the values of the `OS_AS_SIZE` and `OS_DEFAULT_AS_PARTITION_SIZE` environment variables.

You can find more information about environment variables in *Managing ObjectStore*.

Description of `com.odi.applicationName`

Set the `com.odi.applicationName` property to the name of the application for the current client. This allows users of the `LockTimeoutBlocker` class and the `ossvrstat` utility to retrieve information about clients involved in concurrency conflicts. When you set this property, it can provide information about your application to other clients.

Note The `com.odi.applicationName` property is valid only for the first session that you create. This property is ignored by subsequent sessions.

Description of `com.odi.cachedObjectCount`

The `com.odi.cachedObjectCount` property sets the maximum number of objects that can be cached for a transaction. If a transaction has more objects than this property specifies, only the *first* `n` objects are put into the cache. The default value is 10000. The maximum value is `java.lang.Integer.MAX_VALUE`.

This property works together with the `com.odi.cachedObjectTransactionAge` and `com.odi.disableObjectCaching` properties to control automatic object caching. If you do not want automatic object caching enabled, set the `com.odi.disableObjectCaching` property to `true`.

Description of `com.odi.cachedObjectTransactionAge`

The `com.odi.cachedObjectTransactionAge` property specifies the number of transactions that an object is retained in the cache, since the object was last accessed. The range of valid values are 1 to 255. The default value is 255.

This property works with the `com.odi.cachedObjectCount` and the `com.odi.disableObjectCaching` properties to control automatic object caching. If you do not want automatic object caching enabled, set the `com.odi.disableObjectCaching` property to `true`.

Description of com.odi.cacheSize

The `com.odi.cacheSize` property specifies the size of the C++ client cache in bytes. The default is 8 MB. If the value is a `String` that starts with `0x` or `0X`, ObjectStore treats the value as a hexadecimal number. ObjectStore rounds the cache size down to the nearest whole number of pages. You can set a cache size for every session that you create.

You can find information about how to determine the optimum cache size using the `OS_CACHE_SIZE` environment variable in *Managing ObjectStore*.

Description of com.odi.disableObjectCaching

The `com.odi.disableObjectCaching` property allows ObjectStore to decide whether to cache the contents of objects across transactions. When this property is set to `false` (the default), ObjectStore might choose to cache the contents of active objects when they are made hollow. When this property is set to `true`, ObjectStore does not cache the contents of objects across transactions.

When an active object is made hollow when caching is enabled, ObjectStore might choose to cache the object by not calling the `clearContents()` method on it. Objects that are in a hollow state but still have their contents from a prior transaction are considered to be cached. When a cached hollow object is subsequently accessed, ObjectStore checks whether the object's contents are still valid. If the contents are valid, ObjectStore makes the object active without rereading its contents from the database.

Using the `false` setting for this property has the advantage of improving the performance of your application by reducing the time needed to fetch objects. However, the disadvantage is that garbage collection is less effective. You might find that the performance of your application degrades over time as unaccessed objects are being cached across many transactions.

To determine whether your application is running slow because too many objects are being cached, try setting the `com.odi.cachedObjectCount` and `com.odi.cachedObjectTransactionAge` properties to lower values. You can experiment with these values to arrive at a combination of values that is best suited for your application.

You can also decache some objects whose contents are cached, using an iterator and any of the following ObjectStore API methods that manipulate cached objects:

- `getCachedObjects()`

- `getCachedObjectTxnAge()`
- `decache(Object cachedObject)`
- `clearCache()`
- `clearCache(int nTransactions)`

If the number of objects in the heap gets too large, you receive a Java VM heap exception.

The `com.odi.disableObjectCaching` property works with the `com.odi.cachedObjectCount` and `com.odi.cachedObjectTransactionAge` properties to control automatic object caching. If you do not want automatic object caching enabled, set the `com.odi.disableObjectCaching` property to `true`.

Description of `com.odi.disableWeakReferences`

The `com.odi.disableWeakReferences` property defaults to `false`. This means that ObjectStore uses the weak reference facility of the JDK. If you set this property to `true`, it disables the weak reference facility and ObjectStore does not use it.

When you start the first session in a Java process, the setting of the `com.odi.disableWeakReferences` property is in effect for the duration of the Java process. If you terminate the session and start another session with a different value for the `com.odi.disableWeakReferences` property, the new value is ignored.

A weak reference to an object is a reference that does not prevent the object from being garbage collected by the Java VM's garbage collector. ObjectStore uses weak references in its internal object table to hold references to unmodified persistent objects. If your program does not have any references to persistent objects and the reference in the object table is the only reference, the object can be garbage collected. If the persistent object has been modified and the changes have not yet been saved, ObjectStore uses a strong reference, which does not allow the object to be garbage collected.

The weak reference facility in the JDK 1.1 is implemented in such a way that it prevents unmodified persistent objects from being garbage collected. This is corrected in the JDK 1.2. To use your database with the JDK 1.2, you must upgrade it. See *Upgrading Databases for Use with the JDK 1.2* on page 83.

Description of `com.odi.migrateUnexportedStrings`

The `com.odi.migrateUnexportedStrings` property controls what happens when ObjectStore encounters a cross-segment reference to an unexported `String` object. If this property is not set or if it is set to `true`, ObjectStore creates a new `String` object that has the same value as the referenced object. ObjectStore places this new string in the same cluster as the referring object and substitutes this new string for the referenced string. If this property is set to `false`, ObjectStore signals `ObjectNotExportedException` if it encounters a reference to a string in another segment and that string is not exported.

Description of `com.odi.ObjectStoreLibrary`

`com.odi.ObjectStoreLibrary` specifies the name of the native C++ library that contains the ObjectStore schema and native methods for the application.

If you use Java/C++ interoperability, you must specify this property. If you do not use any C++ libraries in your application, you need not specify this property. The standard library supports `com.odi.coll` collections, which are Java peer objects. A custom library is needed for application-specific peer classes.

You must specify a name that is acceptable to `System.loadLibrary()` and not an explicit path. The name should follow platform conventions for library names. If you do not specify this property, ObjectStore uses the standard library, which provides for primary Java objects but not for Java peer objects. If your ObjectStore Java application is accessing C++ classes, you want the library that provides for Java peer objects.

For additional information, see Chapter 4, *Building the Application in Developing Java Applications That Access C++*.

Description of `com.odi.password` and `com.odi.user`

`com.odi.user` and `com.odi.password` allow you to supply a user name and a password when the ObjectStore server has `Name Password` set for the `Authentication Required` server parameter.

Description of `com.odi.product`

`com.odi.product` allows you to run multiple simultaneous sessions against different ObjectStore Java products in the same Java VM. Each ObjectStore Java product runs in its own session.

You must separately obtain each ObjectStore Java product that you want to use. When you purchase ObjectStore, PSE Pro is not included; PSE Pro must be obtained separately.

Your `CLASSPATH` environment variable must include an entry for each ObjectStore Java product that you want to use. For example, if you want to use ObjectStore and PSE Pro, you must have entries for both `osji.jar` and `pro.jar` in your `CLASSPATH`. The order of the entries does not matter.

When you use multiple ObjectStore Java products, they must be compatible with each other. If they are not, `FatalApplicationException` is signaled.

You can set the `com.odi.product` property to one of the following values (case is not significant):

- PSE Pro
- ObjectStore

The `com.odi.product` property applies to one session; it is not global. In other words, each session has a product attribute. After you start a session, you cannot change the value of `com.odi.product` for that session.

With the `com.odi.product` property, a single process can run all the following at the same time:

- Multiple PSE Pro sessions
- Multiple ObjectStore sessions

You can create a session without explicitly setting the `com.odi.product` property. In this case, ObjectStore software checks the Java system property list. If it finds a value for `com.odi.product`, it uses that value. If it does not find a value, the default is that the software looks for PSE Pro, then ObjectStore, and uses the first product it finds.

There are many ways this feature can be useful. For example, an application can

- Open a PSE Pro database in one session and an ObjectStore database in another session and use both
- Copy data from a database created with one product to a database created with another product

The use of `com.odi.product` to copy data among databases created with different products requires several steps. For example, to copy data from a

PSE Pro database to an ObjectStore database, you must do the following in the PSE Pro session:

- 1 Open the source database.
- 2 Read the objects you want to copy.

The `ObjectStore.deepFetch()` method is useful for doing this.

- 3 Commit the transaction with `ObjectStore.RETAIN_READONLY`.
- 4 Close the source database and specify `true` to retain the persistent objects as transient objects.

If you do not close the database, the persistent objects remain associated with the session in which you read them. This prevents another session from storing them in another database.

Then in the ObjectStore session, you must

- 5 Make the objects reachable from the destination database.
- 6 Commit the transaction.

ObjectStore uses transitive persistence to store all reachable objects in the destination database.

Following is an example of code that performs these steps:

```
import com.odi.*;
import com.odi.util.*;
import java.util.Properties;

class CopyToOSJI implements ObjectStoreConstants {
    public static void main(String[] args) {
        /* Create data in a PSE Pro database and then read it out. */
        Properties properties = new Properties();
        properties.put("com.odi.product", "PSE Pro");

        Session.create(null, properties).join();
        Database database = Database.create(
            "pse.odb", ALL_READ | ALL_WRITE);

        Transaction.begin(UPDATE);
        OSVector vector = new OSVector();
        vector.addElement(new Integer(3));
        vector.addElement(new Integer(4));
        database.createRoot("vector", vector);
        Transaction.current().commit();

        Transaction.begin(READONLY);
        vector = (OSVector)database.getRoot("vector");
        ObjectStore.deepFetch(vector);
    }
}
```

```

        Transaction.current().commit(RETAIN_READONLY);

    /* Close the database and specify true
       to retain persistent objects as transient objects. */

    database.close(true);
    Session.leave();

    /* Copy the data to an ObjectStore database. */

    properties.put("com.odi.product", "ObjectStore");
    Session.create(null, properties).join();
    database = Database.create("osji.odb",
        ALL_READ | ALL_WRITE);

    Transaction.begin(UPDATE);
    database.createRoot("vector", vector);
    Transaction.current().commit();
    database.close();
}
}

```

Description of com.odi.stringPoolSize

The `com.odi.stringPoolSize` property allows you to specify the number of newly created strings ObjectStore maintains in the string pool for the current session. The default is "100".

When ObjectStore is about to migrate a string into a segment, it first checks the string pool for an identical string in the same segment. If it finds one, ObjectStore uses the string that is already stored in the segment instead of adding a new identical string to the segment. The information about the strings that are available to be shared is maintained only for the current transaction. The strings that are available to be shared are maintained in a string pool. ObjectStore resets the string pool to empty at the start of each transaction.

For example, suppose you create two instances of a `Person` object in a transaction. In each instance, the value of the `name` field is `Lee`. If you store both instances in the database in the same transaction, and in the same segment, ObjectStore adds only one instance of the string `"Lee"` to the database. This is true even though the Java VM might contain two instances of the string `"Lee"`. When ObjectStore writes the first `"Lee"` string in the segment, ObjectStore notes it in the string pool. Before ObjectStore stores the next instance of `"Lee"` in the segment, it checks the string pool to see if an identical instance is already in the segment. However, if the two `Person` objects were stored in different segments, the two instances of the string

"Lee" would be migrated to the database and stored separately in each segment.

Continuing the example, suppose you use two transactions and you store one instance of `Person` in each transaction. The result is that there are two identical "Lee" strings in the segment. This is because ObjectStore resets the string pool to be empty at the start of each transaction. Consequently, ObjectStore cannot reuse the "Lee" string from the previous transaction.

Caution

If you use `ObjectStore.destroy()` to destroy strings explicitly, you might want to turn off string pooling so you do not inadvertently destroy a string that is shared by different objects. Alternatively, you can use the persistent GC to reclaim strings when they are no longer referenced. Using the GC is usually preferable to explicitly calling `destroy()`, because it is safer to let the persistent GC collect unreachable strings. Also, this approach is often more efficient and results in less database fragmentation.

Description of `com.odi.trapUnregisteredType`

The `com.odi.trapUnregisteredType` property is useful for troubleshooting `ClassCastException`s. The default is that this property is not set and it is usually best to use the default.

When ObjectStore encounters an object of a type for which it does not have information (that is, the type is unregistered), it checks the setting of the `com.odi.trapUnregisteredType` property.

If the property is not set, ObjectStore creates an instance of the `UnregisteredType` class to represent the object of the unknown type. Your application continues to run as long as it does not try to use the `UnregisteredType` object. Often, this can work well because your application has no need for that particular field. However, if you do try to use the object of the unregistered type, ObjectStore signals `ClassCastException`.

If `com.odi.trapUnregisteredType` is set, ObjectStore does not create an `UnregisteredType` object. Instead, it signals `FatalApplicationException` and provides a message that indicates the name of the unregistered class. For additional information, see [Handling Unregistered Types](#) on page 149.

Description of `com.odi.useImmediateStrings`

When you set the `com.odi.useImmediateStrings` property to `true` when a session is created, strings that are eight characters (bytes) or less are stored as immediate values. This means that these strings are not stored as

independent objects in the database, resulting in a smaller database and avoiding the run-time overhead for tracking the Strings in the object table.

When `com.odi.useImmediateStrings` is enabled, Strings that are stored as values of fields in persistent objects are not persistent, unless you explicitly call `ObjectStore.migrate` on the Strings to make them persistent.

By default, the `com.odi.useImmediateStrings` property is disabled (set to `false`). This avoid problems with queries on `com.odi.coll` collection objects. Queries on `com.odi.util` collections are not affected.

Immediate string values are not understood by queries on `com.odi.coll` collection objects. If a query on a `com.odi.coll` collection accesses a String stored as an immediate value, the query interprets the String as a null object. If a class contains indexable fields, immediate string values are not stored in any of the fields for that class. If your query on `com.odi.coll` collection references a class that does not have indexable fields and you have enabled `com.odi.useImmediateStrings`, you should call `ClassInfo.setAllowsPeerQueries` on the `ClassInfo` for that class.

Chapter 4

Managing Databases

You create databases to store your objects. The `Database` class provides the API for creating and managing databases.

The Java interface to ObjectStore supports rawfs databases. For information about rawfs databases, see *Managing ObjectStore*.

Contents

This chapter discusses the following topics:

Creating a Database	50
Creating Segments	53
Creating Clusters	54
Determining Whether a Database, Segment, or Cluster Is Transient	56
Opening and Closing a Database	56
Moving or Copying a Database	61
Performing Garbage Collection in a Database	61
Schema Evolution: Modifying Class Definitions of Objects in a Database	65
Dumping and Loading Databases	73
Destroying a Database	75
Obtaining Information About a Database	75
Grouping Objects in Multiple Clusters and Segments	77
Database Operations and Transactions	81
Upgrading Databases for Use with the JDK 1.2	83

Creating a Database

The `Database` class is an abstract class that represents a database. When you create a database,

- There must be an active session, or `ObjectStore` signals `NoSessionException`.
- A transaction must not be in progress, or `ObjectStore` signals `TransactionInProgressException`.

Databases are cross-platform compatible. You can create databases on any supported platform and access them from any supported platform.

This section discusses the following topics:

- Method Signature for Creating a Database
- Example of Creating a Database
- Result of Creating a Database
- Specifying a Database Name in Creation Method
- When the Database Already Exists
- Installing Schema on Database Creation

Method Signature for Creating a Database

To create a database, call the static `create` method on the `Database` class and specify the database name and an access mode. The method signature is

```
public static Database create(String name, int fileMode)
```

`ObjectStore` signals `AccessViolationException` if the access mode does not provide owner write access.

Example of Creating a Database

For example:

```
import com.odi.*;
class DbTest {
    void test() {
        Database db = Database.create("objectsrus.odb",
            ObjectStore.OWNER_WRITE);
        ...
    }
}
```

This example creates an instance of `Database` and stores a reference to the instance in the variable named `db`. The `Database.create` method is called with two parameters.

The first parameter specifies the pathname of a file. The path can specify a relative name or a fully qualified name. It must always specify a file that is one of the following:

- Local
- In a mounted directory
- In an unmounted remote directory — in this case the file must be identified by a pathname that specifies a remote host

An ObjectStore server must be available for the directory that contains the specified file.

The second parameter specifies the access mode for the database.

Terminology note

`Database` is an abstract class, so ObjectStore actually creates an instance of a subclass that extends `Database`. From your point of view, it does not matter whether ObjectStore creates an instance of `Database` or an instance of a `Database` subclass.

Result of Creating a Database

The result is a database named `"objectsrus.odt"` with an access mode that allows the owner to modify the database. The example stores the reference to the `Database` object in the `db` variable. This means that `db` represents, or is a handle for, the `objectsrus.odt` database.

For each database you create, ObjectStore creates an instance of `Database` to represent your database. Each database is associated with one instance of `Database`. Consequently, you can use the `==` operator to determine whether two `Database` objects in the same session represent the same database. For example, the following method returns `true`:

```
boolean checkIdentity(String dbname) {
    Database db = Database.create(dbname,
        ObjectStore.OWNER_WRITE);
    Database dbAgain = Database.open(dbname,
        ObjectStore.UPDATE);
    return (db == dbAgain);
}
```

Specifying a Database Name in Creation Method

When you create or open a database, you can specify or pass a database name that is a relative name, an absolute operating system pathname, or a rawfs pathname. ObjectStore considers local network mount points when interpreting pathnames. A pathname can refer to a database on a remote host. However, an ObjectStore server must be available to the local host of the directory that contains an ObjectStore database.

If you want to refer to a database on a remote host for which there is no local mount point, you can use a server host prefix. This is the name of the remote host followed by a colon (:), as in `oak:/foo/bar.odb` or `jackhammer:c:\bob\mydb.odb`. In Windows, you can also use UNC pathnames, as in `\\oak\c\foo\bar.odb`.

You can also use locator files to allow access to additional hosts. See *Managing ObjectStore* for information.

When the Database Already Exists

If you try to create a database that already exists, ObjectStore signals `DatabaseAlreadyExistsException`. Before you create a database, you might want to check to see whether it exists and destroy it if it does. For example, you can insert the following before you create a database:

```
try {
    Database.open(dbName, ObjectStore.UPDATE).destroy();
} catch(DatabaseNotFoundException e) {
}
```

Warning

Do this only if you want to destroy and recreate your database. Otherwise, invoke `Database.open()`.

Installing Schema on Database Creation

The `Database.create()` method has an overloading that allows you to install the schema in batch mode rather than incrementally. The default is that ObjectStore performs schema installation as needed. The advantage of batch schema installation is that concurrency conflicts due to schema installation are minimized. The disadvantage is that database creation takes a little longer and the initial database size is larger. See *Installing Schema Information in Batch Mode* on page 271.

Creating Segments

ObjectStore creates each database with one segment that you can use to store objects. A segment is a collection of one or more clusters, which are variable-sized regions of disk space that ObjectStore uses to group objects stored in the database.

The segment is the smallest unit of storage that can be garbage collected, because objects in the same segment can freely reference each other. However, objects in different segments, and even in different databases, can also refer to each other. Such cross-segment references are explicitly tracked in the exported object table. As a result, a single segment can be garbage collected without inspecting the other segments it might refer to or the other segments that might refer to it.

Initially, the size of the segment is about 3 KB, and it consists of a single cluster. As you store additional objects in the segment, ObjectStore increases the size of the segment automatically.

To create additional segments, use the `Database.createSegment()` method. To locate a segment by its ID, use the `Database.getSegment()` method. To retrieve the segment that contains a particular object, call the `Segment.of()` method on that object.

The initial segment in a database is the default segment. To change the default segment, invoke the `Database.setDefaultSegment()` method.

To lock all objects in a segment, see [Locking Objects, Clusters, Segments, and Databases to Ensure Access](#) on page 267.

For additional information about segments, see *Managing ObjectStore*.

Storing Objects in a Particular Segment

You can store objects in different segments by changing the segment that is the default segment or by explicitly migrating an object to a particular segment. See [Chapter 6, Storing, Retrieving, and Updating Objects](#), on page 99, for more information on how to store objects in a specific segment.

To make a persistent reference from an object in one segment to an object in another segment, the object in the other segment (the referenced object) must be exported. See [Grouping Objects in Multiple Clusters and Segments](#) on page 77 for additional information about distributing objects among segments.

Iterating Through the Segments in a Database

To obtain an enumeration of the segments in a database, call the `Database.getSegments()` method. The signature is

```
public DatabaseSegmentEnumeration getSegments()
```

This method returns a `DatabaseSegmentEnumeration` object. After you have this object, you can use the following methods to iterate over the segments in the enumeration:

- `DatabaseSegmentEnumeration.nextElement()`
- `DatabaseSegmentEnumeration.nextSegment()`
- `DatabaseSegmentEnumeration.hasMoreElements()`

If you or another session adds a segment to a database after you create an enumeration, the enumeration might or might not include the new segment. If it is important for the enumeration to list all segments accurately, you should recreate the enumeration after you create the segment.

After you create an enumeration, a segment in the enumeration might be destroyed. If you use the enumeration to try to access a destroyed segment, `ObjectStore` skips the destroyed segment. However, if you retrieve a segment with `DatabaseSegmentEnumeration.nextElement()` or `nextSegment()`, then the segment is destroyed, `ObjectStore` signals `SegmentNotFoundException` if you try to use the destroyed segment.

Creating Clusters

`ObjectStore` creates each segment with one cluster that you can use to store objects. A cluster is a variable-sized region of disk space that `ObjectStore` uses to group objects stored in the database. Initially, the size of the cluster is about 3 KB. As you store additional objects in the cluster, `ObjectStore` increases the size of the cluster automatically.

To create additional clusters in a segment, use the `Segment.createCluster()` method. To locate a cluster by its ID, use the `Segment.getCluster()` method. To retrieve the cluster that contains a particular object, call the `Cluster.of()` method on that object.

Note

For 32-bit platforms, all transient C++ peer objects reside on the same transient cluster. However, for 64-bit platforms, transient C++ peer objects might reside on different transient clusters. Therefore, on 64-bit platforms,

when you call `Cluster.of()` on two different transient C++ peer objects, different clusters might be returned because a different `Cluster` object is returned for each 2^{32} byte range of transient pointers.

The initial cluster in a database is the default cluster. To change the default cluster, invoke the `Segment.setDefaultCluster()` method. To lock all objects in a cluster, see Locking Objects, Clusters, Segments, and Databases to Ensure Access on page 267.

For additional information about clusters, see Managing Databases in Chapter 1 of *Managing ObjectStore*.

Storing Objects in a Particular Cluster

You can store objects in different clusters by changing the cluster that is the default segment or by explicitly migrating an object to a particular cluster. The use of clusters to group objects can improve performance by collocating objects that usually are accessed together and by separating those objects that usually are not accessed together. Distributing two objects into different clusters also eliminates the possibility of spurious lock contention when you access only one of the objects.

Iterating Through the Clusters in a Segment

To obtain an iterator of the clusters in a segment, call the `Segment.getObjects()` method. The signature is

```
public Iterator getObjects()
```

This method returns an iterator object. After you have this object, you can use the following methods to iterate over the segments in the iterator:

```
Iterator.next()
```

```
Iterator.hasNext()
```

If you or another session adds a cluster to a segment after you have created the iterator, the iterator might or might not include the new cluster. If it is important for the iterator to list all clusters accurately, you should recreate the iterator after you have created the new cluster.

After you create an iterator for the objects in a cluster, other sessions are blocked from destroying that cluster until you end your transaction.

If you create the iterator and then destroy the cluster, the next call to `next()` or `hasNext()` causes `ObjectStore` to signal `ClusterNotFoundException`.

There is a bug that makes the `iterator` work incorrectly if you call `Transaction.checkpoint(RETAIN_STALE)`. Doing so causes the next use of the iterator to signal `ObjectException` because of stale objects. This will be fixed in a future release.

Determining Whether a Database, Segment, or Cluster Is Transient

Sometimes there are Java peer objects that identify C++ objects that have been transiently allocated. `ObjectStore` stores these C++ objects in the transient database, transient segment, and transient cluster. To determine whether a database or segment is transient, you can do the following:

```
CPlusPlus.getTransientCluster() == cluster  
CPlusPlus.getTransientSegment() == segment  
CPlusPlus.getTransientDatabase() == database
```

Only Java peer objects are stored in the transient database, segment, or cluster. Java primary objects are never stored in the transient database, transient segment, or transient cluster, even if they are transient.

If you try to retrieve the cluster, segment, or database of a transient primary object, `ObjectStorePSE Pro` signals `ObjectNotPersistentException`.

Opening and Closing a Database

A database can be either open or closed. A database must be open before you can store or access objects in that database. When an application opens a database, it does not matter whether

- A transaction is in progress.
- The database is already open.

When an application closes a database, a transaction cannot be in progress and the database must be open.

This section discusses the following topics:

- Opening a Database
- Possible Open Modes

- Opening the Same Database Multiple Times
- Closing a Database
- Automatic Opens of a Database
- Objects in Closed Databases

Opening a Database

When you open a database, it does not matter whether a transaction is in progress, nor does it matter whether the database is already open.

When you create a database, ObjectStore creates and opens the database. To open an existing database, call the static `Database.open()` method. The method signature is

```
public static Database open(String name, int openMode)
```

For example:

```
Database db = Database.open("myDb.odb",
    ObjectStore.READONLY);
```

The first parameter specifies the pathname of your database. The second parameter indicates the open mode of the database.

Possible Open Modes

ObjectStore provides constants that you can specify for the `openMode` parameter to `Database.open()`. The constants you can specify for `openMode` are

- `ObjectStore.UPDATE` to read and modify a database.
- `ObjectStore.READONLY` to read but not modify a database.
- `ObjectStore.MVCC` allows you to open the database for single-database MVCC (multiversion concurrency control). This allows you to read the database but it does not block another session from updating the database.
- `ObjectStore.MULTI_DB_MVCC` allows you to open the database for multidatabase MVCC. This allows you to specify a set of databases that sessions can read without blocking other sessions from updating those database.

Note

Using `Database.open(name, openMode)` might block the first time you open a database in a session. You can avoid this blocking by opening the database the first time using the `MVCC` mode, closing the database, then using the `Database` object to reopen the database in the mode you want.

Incorrect attempts to modify

If you open a database with `ObjectStore.READONLY`, `ObjectStore.MVCC`, or `ObjectStore.MULTI_DB_MVCC` and attempt to modify an object, `ObjectStore` signals `UpdateReadOnlyException` when you try to commit the transaction.

Example

Suppose you previously created and closed a database that is represented by an instance of a `Database` subclass stored in the `db` variable. You can call the instance `open()` method to open your database in the following way:

```
db.open(ObjectStore.READONLY);
```

You can use the static class `open()` method in the following way:

```
db = Database.open("myDb.odb", ObjectStore.READONLY);
```

Typically, both lines cause the same result. However, they might cause different results if a database has been destroyed and recreated.

To lock all objects in a database, see [Locking Objects, Clusters, Segments, and Databases to Ensure Access](#) on page 267.

Opening the Same Database Multiple Times

After a database is initially opened, if threads in the same session subsequently open the database again, the same database object is returned. For example:

```
db1 = Database.open("foo", ObjectStore.UPDATE);  
db2 = Database.open("foo", ObjectStore.UPDATE);
```

The expression `db1 == db2` returns `true`. They refer to the same database object. Consequently, a call to `db1.close()` or `db2.close()` closes the same database. No matter how many times you open a database, a single call to the `close()` method closes the database. (This is different in the C++ interface to `ObjectStore`. In that interface, for example, if you call `open()` four times and `close()` three times all on the same database, the database is still open.)

Closing a Database

To close a database, call the `close()` method on the instance of the `Database` subclass that represents the database, for example:

```
db.close();
```

The database must be open and there must not be a transaction in progress when `Database.close()` is called; otherwise, an exception is thrown.

Object state after close

When you close a database, all persistent objects that belong to that database become stale or transient. If the last committed transaction that operated on the database retained persistent objects, you can use an overloading of `close()` that allows you to specify what should happen to the retained objects. (For information about retained objects, see [Committing Transactions to Save Modifications](#) on page 122.) The method signature is

```
public void close(boolean retainAsTransient)
```

Specify `true` to make retained objects transient. If you specify `false`, it is the same as calling the `close()` method without an argument. All access to retained objects ends.

Suppose you close a database and make retained objects transient. In the next transaction, if you reread an object from the database that you retained as a transient object, you then have two separate copies of the same object: One is transient and one is persistent. You do not have two references to a single object. When you close a database, all object identity is gone. After you close a database, the database is still associated with the session in which it was closed.

If you do not close

If you do not close a database, `ObjectStore` closes it when you shut down `ObjectStore`.

Database identity

Within a session, `ObjectStore` maintains database identity even after you close a database. For example, consider the following code:

```
import com.odi.*;

public class Goo {
    public static void main(String[] args) {
        Session session = Session.create(null, null);
        session.join();
        try {
            try {
                Database db = Database.create("my.odb", 0664);
                db.close();
            } catch (DatabaseAlreadyExistsException e) {
            }
            Database db1 = Database.open("my.odb",
                ObjectStore.READONLY);
            db1.close();
            Database db2 = Database.open("my.odb",
                ObjectStore.READONLY);
            System.out.println(db1 == db2);
        } finally {
            session.terminate();
        }
    }
}
```

```
}
```

If you run a program with the previous code, the system displays `true`.

In general, it is best to leave databases open for the entire session and write your application so that it shuts down `ObjectStore` before it exits.

However, keeping databases open consumes some `ObjectStore` server resources. Also, there is a limit to the number of databases the `ObjectStore` server can keep open at one time. This depends on the number of open files that the operating system permits, which varies by platform.

Automatic Opens of a Database

Sometimes an application traverses a reference to a database that has not been explicitly opened. `ObjectStore` automatically opens the database according to the default open mode. The default open mode is one of the following:

- `ObjectStore.READONLY`
- `ObjectStore.UPDATE`
- `ObjectStore.MVCC`
- `ObjectStore.MULTI_DB_MVCC`

An application can call the `ObjectStore.setAutoOpenMode()` method to change the default open mode. A call to the `ObjectStore.getAutoOpenMode()` method returns the current open mode. The default autoopen mode is `READONLY`. When you set the autoopen mode, it affects only the current session.

If you know the name of a database that has been opened automatically, you can use `Database.open()` to obtain the already open database. Another way to obtain a handle to an automatically opened database is to call `Database.of()` on an object from the automatically opened database.

If you do not close a database that `ObjectStore` opens automatically, `ObjectStore` closes it when you shut down `ObjectStore`.

Disabling automatic opens

You can disable the ability of `ObjectStore` to automatically open databases. Call the `ObjectStore.setAutoOpenMode()` method and specify the `ObjectStore.DISABLE_AUTO_OPEN` constant. If you do disable automatic opens and your application tries to follow a reference to an unopened database, you receive `DatabaseNotOpenException`.

Objects in Closed Databases

Objects in a closed database are not accessible. However, if you close a database with an argument of `true`, ObjectStore retains the persistent objects as transient objects.

Moving or Copying a Database

You can use the ObjectStore utilities `oscopy` and `osmv` to copy and move a database. You can find information about these utilities in *Managing ObjectStore*.

You can move or copy databases among different supported platforms.

Performing Garbage Collection in a Database

The ObjectStore persistent garbage collector (GC) collects unreferenced Java objects and ObjectStore collections in an ObjectStore database. Persistent garbage collection frees storage associated with objects that are unreachable. It does not move remaining objects to make the free space contiguous.

Contents

This section discusses the following topics:

- Background About the Persistent Garbage Collector
- API for Collecting Garbage in a Database
- API for Collecting Garbage in a Segment
- Command-Line Utility for Collecting Garbage
- Running `osgc` on C++ Databases or Segments

Background About the Persistent Garbage Collector

The ObjectStore persistent GC is independent of the Java VM GC. The Java VM GC is strictly a transient object garbage collector. It never operates on objects in the database.

Applications can continue to use a database while the persistent GC is running. The persistent GC locks portions of a segment as needed, as if it

were just another application. In this way, the GC minimizes the number of pages that are locked and the duration for which the locks are held. The GC also retries operations when it detects lock conflicts.

By default, the GC runs with a transaction priority of zero. Consequently, it is the preferred victim when the server must terminate a transaction to resolve a deadlock. At a later time, the GC redoes the work that was lost when the transaction was aborted.

The GC uses read and write lock time outs of short duration. This avoids competition with other processes for locks. If the GC cannot acquire a lock because of a time-out, it retries the operation at a later time.

The GC performs its job in two major phases. In the mark phase, the GC identifies the unreachable objects. In the sweep phase, the GC frees the storage used by the unreachable objects.

A segment is the smallest storage unit that can be garbage collected. You can specify a segment or a database to be garbage collected. It is usually best to avoid destroying strings (or objects) altogether and let the persistent GC take care of destroying such unreachable objects. The persistent GC typically can destroy and reclaim such objects very efficiently, because it can batch such operations and cluster them effectively. If you set up the GC to run when the system is lightly loaded, you can effectively defer the overhead of the destroy operations to a time when your system would otherwise be idle, thus getting greater real throughput from your application when you really need it.

The persistent GC never removes tombstones for exported objects. For unexported objects, the GC treats tombstones the same way that it treats other objects. The GC removes tombstones if they are not referenced. In other words, the GC removes only unreferenced tombstones. This behavior preserves the safe detection of bad references.

API for Collecting Garbage in a Database

To perform garbage collection on a database, call the `Database.GC()` method. This method invokes the `Segment.GC()` method on each segment in the database. The method signature is

```
public java.util.Properties GC(  
    java.util.Properties GCproperties)
```

For the `GCproperties` parameter, specify null or a `Properties` object for the garbage collection operation. The properties are described in the next section, as they are the same for `Segment.GC()`. If the `GCproperties`

parameter is null, ObjectStore uses the default properties as defined in the documentation for `Segment.GC()`. The properties you can specify are the same as the properties for `Segment.GC()`.

API for Collecting Garbage in a Segment

To perform garbage collection on a segment, call the `Segment.GC()` method. The signature is

```
public java.util.Properties GC(
    java.util.Properties GCproperties)
```

A transaction must not be in progress for the current session. The database that contains the segment you want to garbage collect must not be open for the current session. You cannot perform GC on the transient segment. If you try to, ObjectStore signals `SegmentException`. However, you must open it to create a `Database` object to represent it. After you close the database you want to garbage collect, you can call `Database.GC()` on the `Database` object that represents your closed database.

The `GCproperties` parameter specifies a list of GC properties or null. When you specify null, ObjectStore checks the system properties and uses the default properties, which are suitable for most operations. For more control over GC, you can specify one or more of the following properties. ObjectStore uses the default for any property you do not specify.

- `com.odi.gc.retries` is an `int` that defaults to 10. This indicates the number of times the GC tries to resume the sweep phase of garbage collection after it waits for a lock.
- `com.odi.gc.retryInterval` is an `int` that defaults to 1000. This value indicates the number of milliseconds the sweep operation waits between sweep attempts for a concurrency conflict to be resolved before it tries to resume the sweep.
- `com.odi.gc.lockTimeOut` is an `int` that defaults to 1000. This value indicates the number of milliseconds the sweep operation waits for a lock conflict to be resolved. If it is not resolved in the specified length of time, the GC aborts the current transaction and starts a new transaction. ObjectStore rounds this value up to the nearest second.
- `com.odi.gc.transactionPriority` is an `int` that defaults to 0. This is the transaction priority associated with transactions started by the GC. The server uses this specification when it must determine the transaction that must be the victim in a deadlock. This number is intentionally low so that the GC transaction is the deadlock victim of choice.

- `com.odi.gc.displayGarbage` is an `int` that defaults to 0. If it is not 0, objects that are unreachable are not destroyed. Instead, they are displayed. The argument controls the level of detail in the display:

0	No display
1	Lists the total number of candidates for garbage collection
2	Reserved
3	Lists the location of each candidate for garbage collection
4	Lists the roots of the object graphs of the candidates for garbage collection

The `GC()` method returns a `Properties` object that contains information about the results of the garbage collection. The properties in this object are as follows:

- `com.odi.gc.reclaimedObjects` is the number of objects that were collected by the GC operation.
- `com.odi.gc.reachableObjects` is the number of objects that the GC found to be reachable.

Command-Line Utility for Collecting Garbage

The command-line utility for collecting garbage in a database is `osgc`. See *Managing ObjectStore* for information on how to run the garbage collection utility.

Running `osgc` on C++ Databases or Segments

You can run the `osgc` utility on C++ databases and segments. For more information, see *Managing ObjectStore*.

Performing Compaction on a Database

Your ObjectStore database can become fragmented over time as you create and delete persistent objects. Holes left by the objects that you delete cause this fragmentation. You might want to squeeze out the deleted space in your database to

- Free persistent storage space so it can be used by other clusters in the database.
- Fit more persistent objects on a page. By reducing the number of pages that must be brought into the cache, you can improve the performance of your application.

You can compact your database by using the `oscompact` utility. This utility allows you to specify a database, segment, or cluster for compaction. See *Managing ObjectStore* for information about the `oscompact` utility.

Schema Evolution: Modifying Class Definitions of Objects in a Database

You can modify the class definitions for objects already stored in a database. This process is called schema evolution, because a database schema is a description of the classes whose instances are stored in a database.

There are two primary ways to evolve a schema:

- Use the `Database.evolveSchema()` API.
- Use serialization with a dump/load utility.

You can always use the schema evolution API, but you should use it when the data you must evolve contains very large object graphs. Also, you must use the API when a database contains instances of `com.odi.coll` objects.

Use the serialization technique with the sample code provided only when the database you want to evolve fits into heap space. When you use the serialization technique, the database cannot contain ObjectStore collections because they are not serializable.

The topics discussed in this section are

- When is schema evolution required?
- Preparing to use the schema evolution API
- Using the schema evolution API
- Considerations for using serialization to evolve schema
- Steps for using sample code that uses serialization with a dump/load utility
- Sample code

When Is Schema Evolution Required?

If you change your class in one of the following ways, you must evolve the schema:

- Add or remove a persistent instance field
- Change the type of a persistent instance field (see additional information about indexable fields)
- Change the order of persistent instance fields

ObjectStore initializes each new instance field with its default value as described in section 4.6.4 of the *Java Language Specification*. The new fields are not initialized with the variable initializer, even if the class defines one for the new field.

If you change the type that is associated with a persistent instance field, ObjectStore performs a default initialization, except in cases in which both the old and new instance fields are of the following primitive types: `byte`, `short`, `int`, `float`, or `double`. In this case, ObjectStore applies the appropriate narrowing or widening conversion to the old value and assigns that value to the new instance field. Conversions that involve the `long` type cause ObjectStore to perform a default initialization on the new field.

hashCode() Also, you might need to perform schema evolution if you add or remove the `hashCode()` method. If you use the postprocessor, it determines whether to add a `hashCode()` method. If it previously added a `hashCode()` method and now it does not, or if it previously did not add a `hashCode()` method and now it does, schema evolution is required.

Inheritance You cannot use schema evolution to change the inheritance hierarchy of a class by adding, removing, or changing a superclass.

Allowed changes You can make the following changes to your class and you are not required to evolve the schema:

- Add or remove class or instance methods
- Add or remove class fields
- Add or remove transient instance fields
- Add or remove an implementation of an interface

Indexable fields When you make a field indexable using `com.odi.coll` (peer collections), the change might require schema evolution. If a class, including its superclasses, does not contain any indexable fields, schema evolution is required if you

make a field in the class indexable. The addition of the indexable field changes the representation of the class.

If a class, including its superclasses, has at least one indexable field, schema evolution is not required if you add indexes to other fields.

Preparing to Use the Schema Evolution API

Before you can use the API to perform schema evolution on a database, you must create a `PersistentTypeSummary` instance that identifies the classes whose definitions have changed. The easiest way to do this is to specify the `-summary` option when you run the postprocessor on the updated class definitions. Be sure to run the postprocessor on all the classes in your application at the same time or on a correctly grouped batch of classes.

If the database contains collections that use indexes, you must drop the indexes before you perform schema evolution. After you evolve the schema, you can restore the indexes on the collection.

It is always advisable to make a copy of your database before you evolve its schema. This allows you to restore the database if there are errors.

For you to perform schema evolution, the following conditions must be true:

- The database must not be open for read-only. It can be closed or open for update.
- A transaction must not be in progress.
- There must be a `PersistentTypeSummary` instance that identifies the classes whose definitions have changed.

Using the Schema Evolution API

When ObjectStore performs schema evolution, it makes the database inaccessible to any other operation. To evolve the schema for a database, call the `Database.evolveSchema()` method. The signature is

```
void evolveSchema(String dbName,
                  String workdbName,
                  PersistentTypeSummary summary)
```

The `dbName` parameter specifies the database whose schema you want to evolve. During schema evolution, this database is not available to any other operation.

The `workdbName` parameter specifies the name of a database that ObjectStore uses to hold a checkpoint version of the database while schema evolution

progresses. ObjectStore creates this database as part of schema evolution and destroys it when schema evolution is complete. This working database allows ObjectStore to resume schema evolution if it is interrupted. For example, an interruption can be caused by a power failure or a lack of disk space.

The `summary` parameter specifies a `PersistentTypeSummary` object that identifies the classes whose definitions have changed. It is permissible for the summary to include types that have not changed. However, if the summary does not include a type that has changed, that type might not be evolved.

Typically, you create the summary object as follows:

- 1 When you run the postprocessor on your updated class definitions, specify the `-summary` option.
This creates a class file.
- 2 Call the no-arguments constructor on this class to create the summary object.

Under unusual circumstances, you might create the `PersistentTypeSummary` object yourself. In this case, you must use its constructor for specifying the persistent classes and included summaries.

ObjectStore can evolve only one database at a time. This means that if there are cross-database references, only one database at a time is unavailable to other operations.

When you perform schema evolution on a particular class, you must provide definitions for all superclasses that are part of the schema for the database being evolved.

Considerations for Using Serialization to Perform Schema Evolution

To evolve a schema using serialization, use the following steps:

- 1 Use serialization to dump the contents of a database.
- 2 Modify your class definitions.
- 3 Reload the data.

The next two sections provide instructions for using a sample program that evolves a schema using serialization. Before you use this sample program, you should be aware of the issues described in this section.

Serialize objects	To serialize objects into the database, the classes of all the objects stored in the database must implement <code>java.io.Serializable</code> . If you have a database that contains objects that do not implement <code>Serializable</code> , you can modify the class definitions just to implement <code>Serializable</code> , recompile them, and still access the database. This allows you to dump the database to a file before you make the real class modifications.
<code>readObject()</code>	When you modify a class after doing the dump, you must ensure that the <code>readObject()</code> method considers the old and new versions of the class to be compatible. The most straightforward way to do this is to create a <code>static final long</code> field called <code>serialVersionUID</code> in the modified class. This field must have the same value as the serial version UID for the original class. You can obtain the value for the original class with the <code>serialver</code> utility, for example: <pre>serialver DumpReload\$LinkedList DumpReload\$LinkedList: static final long serialVersionUID = -5286408624542393371L;</pre>
Database size	The database whose schema you want to evolve must be small enough to fit into heap space. If it is not, you must customize the code that dumps and loads the database. You would have to organize your data so that you do not have to serialize all the data in the database at one time.
Large numbers of connected objects	The use of <code>ObjectStore.deepFetch()</code> is a performance concern for very large object graphs. The current implementation of <code>deepFetch()</code> is not careful about bounding stack space. A consequence of this is that it is sometimes impossible to successfully perform the <code>deepFetch()</code> operation for very large object graphs.

Steps for Using Sample Schema Evolution Serialization Code

The sample program that is provided in the next section takes an argument that causes it to perform one of the following three actions:

- Create a database with some data in it, such as instances of `OSHashtable`, `OSVector`, or linked lists
- Use object serialization to dump data in the database to a file
- Use object serialization to reload the data from the file

For example, you can use the sample program to add a new field to the `LinkedList` class. To do so, follow these steps:

- 1 Place the code in a file called `DumpReload.java`.

- 2 Set your `CLASSPATH` environment variable to include the directory that contains the `osjcfpout` file and the `DumpReload.java` file.

- 3 Compile the program with the command

```
javac DumpReload.java
```

- 4 Run the postprocessor to annotate the `DumpReload` and `LinkedList` classes:

```
osjcfp -dest osjcfpout DumpReload.class \  
DumpReload$LinkedList.class
```

- 5 Create the database:

```
java DumpReload create data.odb
```

- 6 Use serialization to dump the data:

```
java DumpReload dump data.odb data.out
```

- 7 Change the `LinkedList` class. Do this by removing the comment flag from the `newField` field in `LinkedList`.

- 8 Recompile the class:

```
javac DumpReload.java
```

- 9 Rerun the postprocessor to annotate the `DumpReload` and `LinkedList` classes:

```
osjcfp -dest osjcfpout DumpReload.class \  
DumpReload$LinkedList.class
```

- 10 Use serialization to reload the data:

```
java DumpReload reload data.odb data.out
```

Sample Code for Using Serialization to Perform Schema Evolution

Following is the sample program that uses serialization to perform schema evolution:

```
import com.odi.*;  
import com.odi.util.OSHashtable;  
import com.odi.util.OSVector;  
  
import java.io.FileInputStream;  
import java.io.FileOutputStream;  
import java.io.IOException;  
import java.io.ObjectInputStream;  
  
import java.io.ObjectOutputStream;  
  
import java.util.Enumeration;
```

```

public class DumpReload {
    static public void main(String argv[]) throws Exception {
        if (argv.length >= 2) {
            if (argv[0].equalsIgnoreCase("create")) {
                createDatabase(argv[1]);
            } else if (argv[0].equalsIgnoreCase("dump")) {
                dumpDatabase(argv[1], argv[2]);
            } else if (argv[0].equalsIgnoreCase("reload")) {
                reloadDatabase(argv[1], argv[2]);
            } else {
                usage();
            }
        } else {
            usage();
        }
    }

    static void usage() {
        System.err.println(
            "Usage: java DumpReload OPERATION ARGS...\n" +
            "Operations:\n" +
            "    create DB\n" +
            "    dump FROMDB TOFILE\n" +
            "    reload TODB FROMFILE\n");
        System.exit(1);
    }

    /* Create a database with 3 roots. Each root contains an
       OSHashtable of OSVectors that contain some Strings. */
    static void createDatabase(String dbName) throws Exception {
        ObjectStore.initialize(null, null);

        try {
            Database.open(dbName, ObjectStore.UPDATE).destroy();
        } catch (DatabaseNotFoundException DNFE) {
        }

        Database db = Database.create(dbName, 0644);

        Transaction t = Transaction.begin(ObjectStore.UPDATE);
        for (int i = 0; i < 3; i++) {
            OSHashtable ht = new OSHashtable();
            for (int j = 0; j < 5; j++) {
                OSVector vec = new OSVector(5);
                for (int k = 0; k < 5; k++)
                    vec.addElement(new LinkedList(i));
                ht.put(new Integer(j), vec);
            }
            db.createRoot("Root" + Integer.toString(i), ht);
        }
        t.commit();
        db.close();
    }
}

```

```
static void dumpDatabase(String dbName, String dumpName)
    throws Exception {
    ObjectStore.initialize(null, null);

    Database db = Database.open(
        dbName, ObjectStore.READONLY);

    FileOutputStream fos = new FileOutputStream(dumpName);
    ObjectOutputStream out = new ObjectOutputStream(fos);

    Transaction t = Transaction.begin(ObjectStore.READONLY);

    /* Count the roots and write out the count. */
    Enumeration roots = db.getRoots();
    int nRoots = 0;
    while (roots.hasMoreElements()) {
        String rootName= (String) roots.nextElement();
        /* Skip internal OSJI header */
        if (!rootName.equals("_DMA_Database_header")) nRoots++;
    }
    out.writeObject(new Integer(nRoots));

    /* Rescan and write out the data.
       The deepFetch() call is necessary because it obtains the
       contents of all objects that are reachable from root,
       and makes them available for serialization. */
    roots = db.getRoots();
    while (roots.hasMoreElements()) {
        String rootName = (String) roots.nextElement();
        if (!rootName.equals("_DMA_Database_header")) {
            out.writeObject(rootName);
            Object root = db.getRoot(rootName);
            ObjectStore.deepFetch(root);
            out.writeObject(root);
        }
    }

    t.commit();

    out.close();
}

static void reloadDatabase(String dbName, String dumpName)
    throws Exception {
    ObjectStore.initialize(null, null);

    try {
        Database.open(
            dbName, ObjectStore.UPDATE).destroy();
    } catch (DatabaseNotFoundException DNFE) {
    }

    Database db = Database.create(dbName, 0644);

    FileInputStream fis = new FileInputStream(dumpName);
    ObjectInputStream in = new ObjectInputStream(fis);

    Transaction t = Transaction.begin(ObjectStore.UPDATE);
```



```

int nRoots = ((Integer) in.readObject()).intValue();
while (nRoots-- > 0) {
    String rootName = (String) in.readObject();
    Object rootValue = in.readObject();
    System.out.println("Creating "+rootName+" "+ rootValue);
    db.createRoot(rootName, rootValue);
}

t.commit();
db.close();
}

static
class LinkedList implements java.io.Serializable {
    private int value;
    private LinkedList next;
    private LinkedList prev;
    //private Object newField;
    static final long serialVersionUID= -5286408624542393371L;

    LinkedList(int value) {
        this.value = value;
        this.next = null;
        this.prev = null;
    }
}
}

```

Dumping and Loading Databases

To upgrade your ObjectStore database from Release 3.0 to 6.0, you must dump and reload your database. ObjectStore has command-line utilities for dumping (`osdump`) and loading (`osload`) that will help to do. For information on how to use these utilities, see *Managing ObjectStore*.

Note

You *cannot* use Java serialization to upgrade your ObjectStore database.

In general, the process you use for dumping and loading a database should follow these steps:

- 1 Run the `osverifydb` utility on the database. If you discover any problems with your database, you should fix them at this time. See *Managing ObjectStore* for information on how to use `osverifydb`.
- 2 Run the `osbackup` utility to back up your current database. See *Managing ObjectStore* for information on how to use `osbackup`.
- 3 Run the `oscopy` utility to make a copy of your current database on which you can test your dump and load procedures.

You should run the `oscopy` utility and not the `copy` command that comes with your operating system. The `oscopy` utility automatically propagates the database transaction log; whereas, the operating system `copy` command does not. See *Managing ObjectStore* for information on how to use `oscopy`.

- 4 Check your database environment. Is everything in the correct path? Check that your ObjectStore libraries, class paths, `include` files, and binary files are set up correctly.
- 5 Test dumping and loading the copy of the database you made with the `oscopy` utility in Step 3.

Testing can help you to determine the amount of time that is needed to successfully dump and load your actual database when it is off line.

When you are done testing, you must schedule a time for dumping and loading your database when your database is off line. No users or applications can access the database during the dump and load process. Usually, this means a database should be dumped or loaded overnight or over the weekend.

If your database needs to be available at all times (7x24), then contact Technical Support for the procedures you must follow to upgrade your database.

- 6 Take your database off line and dump your database by using the `osdump` utility from the current release of the database. For example, if you are upgrading from ObjectStore Release 3.0 to 6.0, then you would dump your database by using the `osdump` utility for ObjectStore Release 3.0.
- 7 Load the output from the `osdump` utility in Step 6 using the `osload` utility from the new release of ObjectStore.

For example, if you are upgrading from ObjectStore Release 3.0 to Release 6.0, then you would load your database by using the `osload` utility from ObjectStore Release 6.0. Note that you do not need to create any intermediate files between the dump and the load procedures.

- 8 Run the `osverifydb` utility on the database after the load completes successfully.
- 9 Run the `oscompact` utility at this time, if your database needs compaction. For more information on the `oscompact` utility, see *Performing Compaction on a Database* on page 64 and *Managing ObjectStore*.
- 10 Run the `osbackup` utility on the upgraded database, before the database goes back on line.

Note

For information about using the preceding procedure to migrate an ObjectStore Release 3.0 database to Release 6.0, see the description of the `osdump` utility in *Managing ObjectStore*.

Destroying a Database

Destroying a database makes all objects in the database permanently inaccessible. You cannot recover a destroyed database except from backups.

To destroy a database, call the `destroy()` method on the `Database` subclass instance, for example:

```
db.destroy();
```

The database must be open for update and a transaction cannot be in progress.

When you destroy a database, all persistent objects that belong to that database become stale.

Obtaining Information About a Database

You can call methods on a database to answer the following questions:

- Is a Database Open?
- What Kind of Access Is Allowed?
- What Is the Pathname of a Database?
- What Is the Size of a Database?
- With Which Session Is the Database or Segment Associated?
- Which Objects Are in the Database?
- Are There Invalid References in the Database?
- Which Databases Are Affiliated with This Database?

Is a Database Open?

To determine whether a database is open, call the `isOpen()` method on the database, for example:

```
db.isOpen();
```

This expression returns `true` if the database is open. It returns `false` if the database is closed or if it was destroyed. To determine whether `false` indicates a closed or destroyed database, try to open the database.

What Kind of Access Is Allowed?

To check what kind of access is allowed for an open database, call the `getOpenMode()` method on the database. The database must be open or `ObjectStore` signals `DatabaseNotOpenException`. The method signature is

```
public int getOpenMode()
```

This method returns one of the following constants:

- `ObjectStore.READONLY`
- `ObjectStore.UPDATE`
- `ObjectStore.MVCC`
- `ObjectStore.MULTI_DB_MVCC`

Following is an example of how you can use this method:

```
void checkUpdate(Database db) {  
    if (db.getOpenMode() != ObjectStore.UPDATE)  
        throw new Error("The database must be open for update.");  
}
```

What Is the Pathname of a Database?

To find out the pathname of a database, call the `getPath()` method on the database, for example:

```
String myString = db.getPath();
```

What Is the Size of a Database?

To obtain the size of a database, call the `getSizeInBytes()` method on the database. The database must be open and a transaction must be in progress, for example:

```
db = Database.open("myDb.odb", ObjectStore.READONLY);  
Transaction tr = Transaction.begin(ObjectStore.READONLY);  
long dbSize = db.getSizeInBytes();
```

This method does not necessarily return the exact number of bytes that the database uses. The value returned might be the result of your operating system's rounding up to a block size. You should be aware of how your operating system handles operations such as these.

With Which Session Is the Database or Segment Associated?

To obtain the session with which a database or segment is associated, call the `Placement.getSession()` method. The method signature is

```
public Session Placement.getSession()
```

Which Objects Are in the Database?

The `osjshowdb` utility displays information about one or more databases. This utility is useful when you want to know how many and what types of objects are in a database. You can use this utility to verify the general contents of the database.

Information about the `osjshowdb` utility is in `osjshowdb: Displaying Information About a Database` on page 339.

Are There Invalid References in the Database?

The `osjcheckdb` utility or the `Database.check()` method checks the references in a database. This tool scans a database and checks that there are no references to destroyed objects. The most likely cause of dangling references is an incorrectly written program. You can fix dangling references by finding the objects that contain them and overwriting the invalid references with something else, such as a null value. In addition to finding references to destroyed objects, the tool performs various consistency checks on the database.

Information about the `osjcheckdb` utility is in `osjcheckdb: Checking References in a Database` on page 338.

Which Databases Are Affiliated with This Database?

The `Database.getAffiliatedDatabases()` method returns an iterator of those databases associated with this database. Databases whose objects can be referenced are known as *affiliated databases*. The iterator can only be used within the transaction in which it is created.

See `Referencing Objects Stored in Other Databases` on page 117 for information on how to affiliate databases.

Grouping Objects in Multiple Clusters and

Segments

You can improve the performance of your application by grouping together objects that are expected to be used together. Effective grouping does two things:

- Reduces the number of disk and network transfers an application requires
- Increases concurrency among applications

Use clusters to group objects

You can use clusters to store groups of objects that are accessed together. Objects in one cluster can refer to objects in any other cluster in the same segment without incurring performance overhead.

Objects can also refer to other objects in different segments or different databases using cross-segment references and cross-database references. However, using these types of references require that the referenced object be explicitly exported. Exported objects incur some performance overhead.

One segment and multiple clusters

It generally is easier and more efficient to have only one segment and to group objects using multiple clusters in that segment to avoid exporting objects.

Garbage collecting and segments

There are times when you might want more than one segment, especially for garbage collection. The segment is the smallest storage unit that can be garbage collected; therefore, it might be desirable to break large databases into more than one segment for purposes of garbage collecting.

When there is more than one segment in a database, you must carefully plan for any cross-segment references you might use. The following section describes a procedure you should follow if you intend to have more than one segment in your database.

Planning for Cross-Cluster and Cross-Segment References

To obtain the best performance possible from your application, it is important that you group your objects appropriately into clusters and segments. As a rule, you should avoid placing objects in more than one segment. To determine where and how to place your objects, follow these steps:

Create a plan

- 1 Create a plan for organizing your objects into clusters and for assigning the clusters to segments. You might want many clusters in a single segment or a small number of segments.

Base your plan on the size of the objects and the patterns of access, grouping objects that tend to be accessed together in the same cluster.

Try to use as few segments as possible, unless it is necessary to garbage collect the database in small (segment) increments.

Placing objects in one segment

- 2 Most objects probably will become persistent by reachability. This means that an object is stored in the same segment and cluster as the object that refers to it. If an object becomes reachable because it is a value of a database root, it is stored in the default segment and cluster.

To begin placing your objects in clusters, use the `ObjectStore.migrate()` method to explicitly store a few starter objects in specific segments and clusters. Then, subsequently created objects that are reachable by these starter objects are automatically migrated to the same cluster.

Placing objects in multiple segments

- 3 If you need to place objects in more than one segment for garbage collecting purposes and these objects need to access objects in the other segments, you need cross-segment references.

To create cross-segment references for each of the planned segments, you must designate a few target objects to be the points of reference for objects from the other segments. These target objects are usually the same objects as the starter objects in step 2.

Once again, you use the `ObjectStore.migrate()` method to store the designated (target) objects in their respective clusters and segments, but this time you must specify `true` for the `Export` argument.

Placing objects in multiple databases

- 4 If you need to place objects in multiple databases, see [Referencing Objects Stored in Other Databases](#) on page 117 for information on creating cross-database references.

Caution

If you are using cross-segment and cross-database references, you must ensure that they are referring to exported objects only. An object is only exported if any of the following is true:

- It is the value of a database root.
- The `ObjectStore.export()` method was called on it.
- It was stored using the `ObjectStore.migrate()` method with `true` as the `Export` argument.

Minimize exported objects

You want to minimize the number of exported objects in each segment. A segment that contains a large number of exported objects has more overhead because of the extra data needed to describe the objects and the extra time it takes for an application to reference them.

You need to call `ObjectStore.migrate()` only on objects that are not already stored in the database. After objects are stored in the database, use the `ObjectStore.export()` method to allow objects in other segments to reference them. After an object is stored in a database, you cannot change its location.

Objects that are referenced by database roots are exported automatically.

Export Exceptions

Objects can become persistent by reachability. When this happens, `ObjectStore` places the newly persistent object in the same segment as the object from which it can be reached. If objects from other segments are also referencing the newly persistent object, it must be explicitly migrated or exported.

If you create a situation in which an object is referred to by objects from other segments and the object is not explicitly exported, `ObjectStore` signals `ObjectNotExportedException`. If the object is referred to by objects in a nonaffiliated database, `ObjectStore` signals `DatabaseNotAffiliatedException`.

Because `ObjectStore` does not trace every storage operation, it does not check whether an object has been explicitly exported when you store the reference. Instead, the check occurs and the `ObjectNotExportedException` is signaled when you try to commit the transaction or evict the object that references the unexported object.

Exported objects are not retained on a commit, checkpoint, or abort when the `ObjectStore.RETAIN_STALE` argument is specified. Instead, use external references or the `ObjectStore.RETAIN_HOLLOW` argument value.

With the exception of `ObjectStore` collection objects, you cannot export peer objects, only primary objects. If you try to export a peer object that is not an `ObjectStore` collection object, `ObjectStore` signals `ObjectException`.


How Many Exported Objects Are Needed?

You do not want too many objects in a segment to be referred to by objects in other segments. In other words, you want to minimize the number of

exported objects in each segment. However, there is no additional cost if you have many objects that refer to the same exported object.

The overhead of having many exported objects in a segment is the result of having many entries in the export table. Access to an exported object is through the export table. Having many exported objects is less efficient than having clustered objects because access to the exported objects requires an indirect look-up through the export table. While the look-up is implemented with a highly tuned hashing mechanism, the overhead for maintaining the export table increases as it grows. Finally, if the exported objects are very volatile, the database locking mechanism can result in additional wait time for users who are creating or accessing exported objects.

For these reasons, it is desirable to design your application so that only a relatively small number of nonvolatile objects must be exported. The following figure shows the optimum way to set up cross-segment references:



Database Operations and Transactions

For each database operation, there are rules about whether it can be performed

- Inside a transaction
- Outside a transaction
- Both inside and outside a transaction

The following table shows the rules that apply to certain operations. If your application tries to perform a database operation that breaks a rule, you receive a run-time exception.

<i>Database Operation</i>	<i>Can Be Performed Inside/Outside Transaction?</i>	<i>Database Open?</i>
<code>acquireLock()</code>	Inside	Open
<code>affiliate()</code>	Both	Open
<code>check()</code>	Inside	Open
<code>close()</code>	Outside	Open
<code>create()</code>	Outside	Not applicable
<code>createRoot()</code>	Inside	Open
<code>createSegment()</code>	Inside	Open
<code>destroy()</code>	Outside	Open
<code>destroyRoot()</code>	Inside	Open
<code>GC()</code>	Outside	Closed
<code>getAffiliatedDatabases()</code>	Inside	Open
<code>getDefaultSegment()</code>	Inside	Open
<code>getOpenMode()</code>	Both	Open
<code>getPath()</code>	Both	Open
<code>getRoot()</code>	Inside	Open
<code>getRoots()</code>	Inside	Open
<code>getSegment()</code>	Inside	Open
<code>getSegments()</code>	Inside	Open
<code>getSizeInBytes()</code>	Inside	Open
<code>installTypes</code>	Inside	Open
<code>isOpen()</code>	Both	Open or closed
<code>of()</code>	Inside	Open
<code>open()</code>	Both	Open or closed
<code>setDefaultSegment()</code>	Inside	Open

<i>Database Operation</i>	<i>Can Be Performed Inside/Outside Transaction?</i>	<i>Database Open?</i>
<code>setRoot()</code>	Inside	Open
<code>show()</code>	Inside	Open

Upgrading Databases for Use with the JDK 1.2

The JDK 1.2 computes hash codes for `String` types differently from the JDK 1.1. As a result, databases that depend on `String` hash codes force the database to be usable only from the JDK 1.1 or the JDK 1.2, but not both.

`ObjectStore` marks databases to specify whether the JDK 1.1 or JDK 1.2 was used when the database was created. It then ensures that the same JDK version is used when the database is accessed.

Databases created with releases before this one are assumed to have been created with the JDK 1.1. With this release of `ObjectStore`, you can use only the JDK 1.1 to access these databases unless you upgrade them for use with the JDK 1.2. `ObjectStore` provides a tool and an API for

- Upgrading databases created with the JDK 1.1 to be accessible with the JDK 1.2
- Marking databases as already created with the JDK 1.2

After you upgrade a database, you can no longer use it with the JDK 1.1.

To use the upgrade tool, see `osjuphsh`: Upgrading String Hash Codes in Databases on page 340. To use the API, see `com.odi.Upgrade.upgradeDatabaseStringHash()` in the *Java API Reference*.

The upgrade facility also allows you to mark a database as not containing any objects that depend on `String` hash codes. If you mark a database in this way, you can use either the JDK 1.1 or the JDK 1.2 to access the database.

After you upgrade your database, Technical Support recommends that you run the GC on the database. Upgrading the database creates some garbage.

If you use the upgrade API to perform the upgrade, watch out for a problem that involves cross-database references from the upgraded database to other databases. If the database being upgraded contains references to other databases, those other databases might need to be opened during the upgrade. If they are opened, the upgrade API leaves them open when it is done. The upgrade allows the other databases to be opened regardless of whether they have been upgraded, because of the limited way in which the other databases are required by the upgrade. However, if the application tries to use those databases after performing the upgrade, it receives exceptions if the other databases are not already upgraded.

The work around is to ensure that you upgrade all databases you intend to use before you try to access any upgraded database.

Chapter 5

Working with Transactions

A transaction is a logical unit of work that runs in a session. Only one transaction can be running in a session at a time. A transaction is a consistent and reliable portion of the execution of a program.

In your code, you place calls to the ObjectStore API to mark the beginnings and ends of transactions. Initial access to a persistent object must always take place inside a transaction. Depending on how the transaction is committed, additional access to persistent objects might be possible.

Either the database is updated with all of a transaction's changes to persistent objects or the database is not updated at all. If a failure occurs in the middle of a transaction, or you decide to abort the transaction, the contents of the database remain unchanged.

A transaction can obtain a write lock or a read-only lock on the database opened by the session. A write lock prevents other sessions in the Java VM from writing to the database. A read-only lock allows other sessions to access the database by using read-only transactions.

Contents

This chapter discusses the following topics:

Starting a Transaction	86
Working Inside a Transaction	87
Ending a Transaction	89
Handling Automatic Transaction Aborts	92
Determining Transaction Boundaries	96

Starting a Transaction

ObjectStore provides the `com.odi.Transaction` class to represent a transaction. You should not make subclasses of this class.

This section discusses the following topics:

- Calling the `begin()` Method
- Allowing Objects to Be Modified in a Transaction
- Difference Between Update and Read-Only Transactions

Calling the `begin()` Method

To start a transaction, call the `begin()` method on the `Transaction` class. This returns an instance of `Transaction` and you can assign it to a variable. The method signature is

Method
signature

```
public static Transaction begin(int type)
```

The transaction type determines whether ObjectStore waits for a write lock on a page. There can be only one write lock on a database, but there can be multiple read-only locks on a database. The type of the transaction can be `ObjectStore.UPDATE` or `ObjectStore.READONLY`.

If there is no open database when you start the current transaction, ObjectStore tries to obtain a read-only lock or a write lock as soon as the session tries to open a database, depending on the type of transaction.

Example

```
Transaction tr = Transaction.begin(ObjectStore.UPDATE);
```

This example returns a `Transaction` object that represents the transaction just started. The result is stored in `tr`. This is an update transaction, which means that the application can modify database contents.

Allowing Objects to Be Modified in a Transaction

To modify persistent objects, you must specify the transaction type to be `ObjectStore.UPDATE`. Also, any database you modify must have been opened for update. Note that even if you open a database for read-only, ObjectStore allows you to start an update transaction. An application does not receive an exception until it tries to modify persistent objects inside the read-only database.

If you try to modify persistent data in a read-only transaction, ObjectStore signals `UpdateReadOnlyException`.

Difference Between Update and Read-Only Transactions

You can start a transaction for `READONLY` or for `UPDATE`. The major difference between the two transaction types is that when you start a transaction for `READONLY`, ObjectStore performs additional checks during the transaction and when you commit the transaction. These checks ensure that changes are not saved in the database if they were made in a read-only transaction. There is no difference in performance between a read-only transaction and an update transaction.

Working Inside a Transaction

A transaction is associated with the session that is associated with the thread that starts the transaction. A transaction remains active until you explicitly commit it or until it aborts. A session can have only one active transaction. Concurrent transactions must be in separate sessions.

This section discusses the following topics:

- Obtaining the Session Associated with the Current Transaction
- Transaction Already in Progress
- Obtaining Transaction Objects
- Performing a Transaction Checkpoint
- Setting a Transaction Priority

Separate transactions that access the same database compete with one another for locks on the objects that they access. This can cause one transaction to wait for another to release its locks. Alternatively, it can cause a transaction deadlock situation in which two or more transactions wait for each other. This forces one transaction to abort.

Two transactions can never update the same object at the same time. However, two transactions can both open the same database for update at the same time and they can make updates concurrently to different parts of the database.

Obtaining the Session Associated with the Current

Transaction

The current session is the session that a thread most recently joined. To obtain the session that is associated with the current transaction, call the `Transaction.getSession()` method. The method signature is

```
public Session Transaction.getSession()
```

To obtain the transaction that is associated with the current session, call the `Session.currentTransaction()` method. The method signature is

```
public Transaction Session.currentTransaction()
```

To determine whether there is a transaction in progress for the current session, call the `Transaction.inTransaction()` method on `Transaction`. The method signature is

```
public static boolean inTransaction()
```

This method returns `true` if there is a transaction in progress for the current session. Otherwise, it returns `false`. It is worth noting that `inTransaction()` returns `false` if the calling thread is not joined to the current session. This can be important if you use an unassociated thread to check whether there is a transaction, and then try to close the database based on a false response. The previously unassociated thread would automatically be joined to the session to close the database. If a transaction is actually in progress, `ObjectStore` signals `TransactionInProgressException`, which is, of course, unexpected because `inTransaction()` returned `false`.

Transaction Already in Progress

Nested transactions are not allowed. If you try to start a transaction when a transaction for the current session is already in progress, `ObjectStore` signals `TransactionInProgressException`.

Obtaining Transaction Objects

An application can obtain the transaction object for the current thread by calling the static `current()` method on the `Transaction` class. The method signature is

```
public static Transaction current()
```

This method returns the transaction object associated with the current session, for example:

```
Transaction.current().commit()
```


This example commits the current transaction. If no transaction is in progress, `current()` signals `NoTransactionInProgressException`.

Performing a Transaction Checkpoint

You can use the `Transaction.checkpoint()` method to commit changes and still continue working with the same persistent objects if you used the `RETAIN_HOLLOW` or `RETAIN_READONLY` argument. When you call the `checkpoint()` method, you specify whether to retain persistent objects as hollow objects or to make all persistent objects stale. This is useful when you are trying to improve concurrency or when you are at a consistent state and want to save your changes but continue working. Any database lock that you hold will continue to be held through the checkpoint. See [Checkpoint: Committing and Continuing a Transaction](#) on page 264 for more information. This means that unlike committing a transaction, you do not lose your database lock, so the work accomplished before the checkpoint remains valid. After the checkpoint, if there is a call to `abort()`, `ObjectStore` rolls back the state of the database to the last checkpoint.

The method signature is

```
public void checkpoint(int retain)
```

For information on the different `retain` values that can be passed to the `checkpoint()` method, see [Committing Transactions to Save Modifications](#) on page 122.

Setting a Transaction Priority

When there is a deadlock, the server uses the transaction priority as one of the criteria to determine the transaction to abort. See [Helping Determine the Transaction Victim in a Deadlock](#) on page 271.

Ending a Transaction

When transactions terminate successfully, they commit and their changes to persistent objects are saved in the database. When transactions terminate unsuccessfully, they abort and their changes to persistent objects are discarded.

For read-only transactions, there are no advantages to committing them rather than aborting them, nor to aborting them rather than committing them.

This section discusses the following topics:

- Committing Transactions
- What Can Cause a Transaction Commit to Fail?
- Aborting Transactions

Committing Transactions

ObjectStore provides the `Transaction.commit()` method for ending a transaction successfully. When an application commits a transaction, ObjectStore

- Saves and commits any changes in the database
- Performs transitive persistence if applicable (see page 11)
- Sets the state of persistent objects that were accessed or referenced in the transaction

Transitive persistence

When ObjectStore commits a transaction, it checks to see if there are any transient objects that are referred to by persistent objects. If there are, and if all referred-to objects are persistence-capable objects, ObjectStore stores the referred-to objects in the cluster that contains the referring object. This is the process of transitive persistence. If any referred-to object is not persistence capable, ObjectStore signals `ObjectNotPersistenceCapableException`.

If a transient object is referred to by multiple objects in different clusters in the same segment, the transient object is migrated to the cluster of the first such referring object encountered during the commit process. However, if a transient object is referred to by multiple objects in different segments, ObjectStore signals `AbortException`.

All cross-segment references must be to exported objects, which means you must have explicitly migrated the object to one of the segments and specified that it is exported.

Making objects stale

To commit a transaction and make the state of persistent objects stale, call the `commit()` method with no argument. For example:

```
tr.commit();
```

The method signature is

```
public void commit()
```

Setting object state

To commit a transaction and be flexible about the state of persistent objects after the transaction, call the `commit(retain)` method on the transaction.

The values you can specify for `retain` are described in [Committing Transactions to Save Modifications](#) on page 122. The method signature is

```
public void commit(int retain)
```

The following example commits the transaction and specifies that the contents of the active persistent objects should remain available to be read.

```
tr.commit(ObjectStore.RETAIN_READONLY);
```

What Can Cause a Transaction Commit to Fail?

When `ObjectStore` tries to commit a transaction, if `ObjectStore` encounters any of the situations in the following list, it causes the transaction commit to fail. When `ObjectStore` aborts a transaction commit, it signals `AbortException`.

- A nonexported object is reachable from an object in a different segment.
- A persistent object references an object that is not persistence capable.
- A persistent object was updated to reference a stale object.
- A deadlock was encountered during commit. This is more likely to happen when lazy write locking is enabled.
- There is a server error. This can happen when there is not enough disk space to fit everything you stored in the database. It can also happen because of a broken network connection, a server failover, or a disk error, making it impossible for the server to complete the commit. Failover also causes `ObjectStore` to abort the transaction.

Aborting Transactions

`ObjectStore` provides the `Transaction.abort()` method for ending a transaction unsuccessfully. An abort can happen explicitly through the `Transaction.abort()` method or implicitly because a session is terminated or there is a system exception. When an application aborts a transaction, `ObjectStore`

- Ensures that the objects in the database are as they were just before the aborted transaction started
- Sets the state of persistent objects from the transaction
- Returns any transient objects that were made persistent during the transaction to their transient state

Transient objects	<p>Only the state of the database is rolled back. The state of transient objects is not undone automatically. For example, if you created new transient objects during the transaction, they still exist after the transaction aborts. Applications are responsible for undoing the states of transient objects. Any form of output that occurred before the abort cannot be undone.</p>
Open databases	<p>If you opened any databases during the transaction, <code>ObjectStore</code> keeps them open. Any databases that were open before the aborted transaction was started remain open after the abort operation.</p>
Application failure	<p>If an application fails during a transaction, when you restart the application the database is as it was before the transaction started. If an application fails during a transaction commit, when you restart the application, either the database is as it was before the transaction that was being committed or the database reflects all the transaction's changes. This depends on how far along in the commit process the application was when it terminated. Either all or none of the transaction's changes are in the database.</p>
<code>abort()</code>	<p>To abort a transaction and set the state of persistent objects to the state specified by <code>Transaction.setDefaultAbortRetain()</code>, call the <code>abort()</code> method. The default state is <code>stale</code>. The method signature is</p> <pre>public void abort()</pre> <p>For example:</p> <pre>tr.abort();</pre>
<code>abort(retain)</code>	<p>To abort a transaction and specify a particular state for persistent objects after the transaction, call the <code>abort(retain)</code> method on the transaction. The values you can specify for <code>retain</code> are described in Specifying a Particular State for Persistent Objects on page 140. The method signature is</p> <pre>public void abort(int retain)</pre> <p>The following example aborts the transaction and specifies that the contents of the active persistent objects should remain available to be read.</p> <pre>tr.abort(ObjectStore.RETAIN_READONLY);</pre>

Handling Automatic Transaction Aborts

`ObjectStore` sometimes aborts a transaction automatically because

- There is a transaction deadlock.

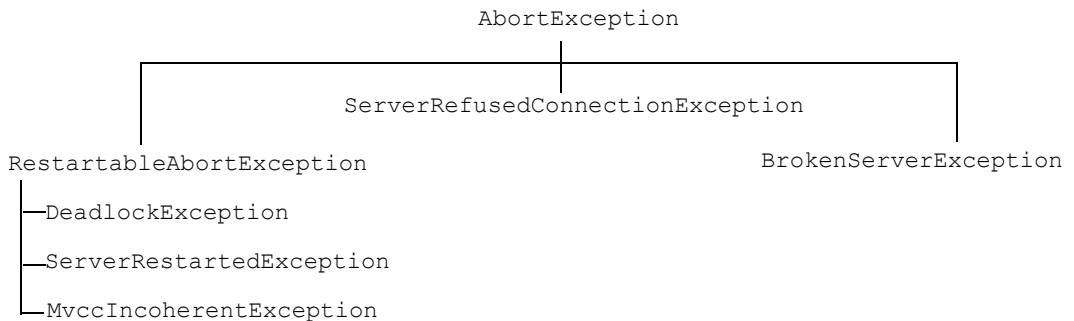
- The connection to the server is broken.
- The server has been restarted.
- The server refuses the connection.
- An exception occurs during a transaction commit. For example, ObjectStore aborts a transaction if you try to commit it when a persistent object refers to a transient object that is not persistence capable. This type of automatic abort is discussed in various sections of Chapter 6, Storing, Retrieving, and Updating Objects, on page 99.

Results of Transaction Abort

When ObjectStore aborts a transaction, it rolls back the persistent state to what it was before the transaction. ObjectStore does not roll back the transient state. Any form of output that occurred before the abort cannot be undone. Therefore, it is generally good practice to perform output outside a transaction.

Description of Transaction Abort Exceptions

When ObjectStore aborts a transaction, it signals an exception that indicates the reason for the abort and whether it makes sense to retry the transaction. Following is the hierarchy of transaction abort exceptions:



The superclass of exceptions for transaction abort is `AbortException`. The superclass of exceptions for which it makes sense to retry the aborted exception is `RestartableAbortException`. `AbortExceptions` that do not extend `RestartableAbortException` indicate that before the transaction can proceed, some action is probably required to correct the problem that caused the exception.

Exceptions that require intervention

Exceptions for which some action is probably required include

- `AbortException` — Signaled during transaction commit.
- `ServerRefusedConnectionException` — Signaled when the server is unavailable when the client tries to connect to it.
- `BrokenServerException` — Signaled when the server becomes unavailable while the client is connected to it. It indicates that the server refused to continue the connection.

If your application receives these exceptions, you should check the situation before you retry the transaction. You should not handle these exceptions with retry loops.

Exceptions that indicate retrying

A `RestartableAbortException` indicates that it makes sense to retry the aborted transaction immediately without taking any other action. This class has three subclasses that indicate the particular reason for the abort.

- `DeadlockException` indicates that transactions were deadlocked. It is possible that the next transaction will also deadlock. However, there is a reasonable chance that the transaction that was allowed to proceed has made progress so that the deadlock is not repeated.
- `ServerRestartedException` indicates that the server was restarted, for example, when failover occurs. The client was already connected to the server. The client receives this exception when it next tries to contact the server. `ObjectStore` aborts the transaction because it cannot commit it. However, things should be fine for a subsequent transaction. If your application receives this exception, it should roll back transient state before it retries the transaction.
- `MvccIncoherentException` indicates that `ObjectStore` can no longer guarantee that all the databases open for multidatabase MVCC will have the same snapshot of data, that is, they are no longer coherent. This can happen, for example, if, by following a cross-database pointer, an attempt is made to open a database in multidatabase MVCC mode within an ongoing transaction. Retrying the transaction will re-establish the coherency of the set of databases opened for multidatabase MVCC, including the one that caused the exception to be thrown.

Restarting Aborted Transactions

`ObjectStore` does not retry transactions that are aborted automatically, even if they extend `RestartableAbortException`. If you want to retry the transaction when your application receives a `RestartableAbortException`,

you must include code to do this in your application. An example of this code follows. If you run the same example in separate VMs at the same time, it produces deadlocks.

```
import com.odi.*;
class DeadlockTest {
    static final int MAX_RETRIES = 10;

    public static void main(String[] args) {
        ObjectStore.initialize(null, null);
        Database db = getDatabase();
        while (true) {
            test(db);
        }
    }

    static void test(Database db) {
        int retries;
        for (retries = 0; retries < MAX_RETRIES; retries++) {
            try {
                Transaction.begin(ObjectStore.UPDATE);
                Integer value = increment(db);
                Transaction.current().commit();
                System.out.println("Value = " + value + ", retries = " +
                    retries);
                break;
            } catch (RestartableAbortException e) {
            }
        }
        if (retries >= MAX_RETRIES)
            System.out.println("Gave up after " + retries +
                " retries.");
    }

    static Database getDatabase() {
        try {
            return Database.open("test.odt", ObjectStore.UPDATE);
        } catch (DatabaseNotFoundException e) {
            Database db = Database.create("test.odt", 0664);
            Transaction.begin(ObjectStore.UPDATE);
            db.createRoot("root", null);
            Transaction.current().commit();
            return db;
        }
    }

    static Integer increment(Database db) {
        Integer value = (Integer)db.getRoot("root");
        if (value == null)
            value = new Integer(0);
        else
            value = new Integer(value.intValue() + 1);
    }
}
```

```
        db.setRoot("root", value);  
        return value;  
    }  
}
```

Handling Deadlocks

A simple deadlock occurs when one transaction holds a lock on a page that another transaction is waiting to access, while at the same time, this other transaction holds a lock on a page that the first transaction is waiting to access. Neither process can proceed until the other does. There are other, more complicated forms of deadlock that are analogous.

ObjectStore has a deadlock detection facility that breaks deadlocks, when detected, by aborting one of the transactions involved in the deadlock. By aborting one transaction (the victim), ObjectStore causes its locks to be released so other processes can proceed.

ObjectStore signals `DeadlockException` when it detects a deadlock. This causes ObjectStore to abort the transaction that causes the deadlock. You can change the transaction that ObjectStore aborts by changing the setting of the `Deadlock Victim` server parameter.

ObjectStore does not detect deadlocks when the deadlock is distributed across multiple servers.

Determining Transaction Boundaries

When determining whether to commit a transaction, consider database state, whether to combine transactions, and interdependencies among cooperating threads.

Inconsistent Database State

You should not commit a transaction if the database is in a logically inconsistent state. A database is considered to be in an inconsistent state if at that moment a just-started transaction would encounter problems on viewing the current state of the data.

Consider your database to be something that moves from one consistent state to another. You should commit a transaction only when the state is consistent. When is a database consistent? If you start your application at this very moment, is the database completely usable exactly the way it is now?

For example, suppose your database contains information about married couples. Couples refer to one another through a `spouse` field. At a particular moment, suppose a person in the database refers to another person in the database through its `spouse` field, but that spouse does not refer to the first person. At that moment, the database is in an inconsistent state.

Another inconsistency you should avoid is migrating objects into a database that are not reachable from any root. Such objects become unreachable if the application fails between transactions or they are removed inappropriately from the database when the garbage collector runs.

When the database state is consistent, you might decide not to commit the transaction. However, if you do not commit, you risk losing changes if `ObjectStore` aborts the transaction. You should always commit changes before you inform a user or another interface that a particular task was accomplished.

Combining Transactions

The transaction commit operation requires network interaction with the `ObjectStore` server. Consequently, you might be able to improve performance by combining several logical transactions into a single transaction.

To do this, skip the call to `commit()` and subsequent call to `Transaction.begin()` for a series of sequential transactions. If many of the same objects are used in the different transactions, the single combined commit operation is more efficient than many small commit operations would be.

However, this strategy means that you risk losing changes if `ObjectStore` aborts the transaction before you commit it. All the logical transactions up to that point would be rolled back with the current logical transaction. You should always commit changes before you inform a user or some other interface that a particular task was accomplished.

Multiple Cooperating Threads

If your application uses cooperating threads, you must take this into account when determining when to commit transactions. For example, you do not want to create a situation in which one thread commits a transaction while a cooperating thread is updating persistent objects. The `commit()` method might make all persistent objects stale for all cooperating threads. If the `commit()` method retains persistent objects, `ObjectStore` discards any

modifications to retained persistent objects at the start of the next transaction. You must coordinate the `Transaction.begin()` and `Transaction.commit()` operations among cooperating threads.

Synchronizing threads is like having a joint checking account. Suppose the amount in the checking account is \$100.00. Your partner writes a check for \$50.00. Then you try to cash a check for \$75.00. This does not work. It does not matter that it was your partner and not you who wrote the check for \$50.00. You and your partner have to cooperate.

Performance Considerations

Committing a transaction, even a read-only transaction, has a certain amount of overhead associated with it. If you have a lot of small transactions, you might want to combine some of them into larger transactions.

Chapter 6

Storing, Retrieving, and Updating Objects

This chapter provides information about how to store data in a database and how to read it back and update it. An application can access persistent data only inside a transaction and only when the database is open.

Contents

This chapter discusses the following topics:

Storing Objects	100
Retrieving Persistent Objects	102
Working with Database Roots	104
Iterating Through the Objects in a Cluster, Segment, or Database	108
Using External References to Stored Objects	109
Referencing Objects Stored in Other Databases	117
Updating Objects in the Database	118
Committing Transactions to Save Modifications	122
Evicting Objects to Save Modifications	132
Aborting Transactions to Cancel Changes	138
Destroying Objects in the Database	142
Default Effects of Various Methods on Object State	146
Transient Fields in Persistence-Capable Classes	146
Avoiding <code>finalize()</code> Methods	148
Troubleshooting Access to Persistent Objects	148
Handling Unregistered Types	149
Troubleshooting <code>OutOfMemoryError</code>	154

Storing Objects

ObjectStore's Java API preserves the automatic storage management semantics of Java. Objects become persistent when they are referenced by other persistent objects. This is called persistence by reachability or transitive persistence. The application defines persistent roots and, when it commits a transaction, ObjectStore finds all objects reachable from persistent roots and stores them in the database.

To store objects in a database, do the following:

- 1 Open the database or create the database in which you want to store the objects. Be sure the database is opened for update. See page 50.
- 2 Start an update transaction. See page 86.
- 3 Create a database root or access an existing database root and specify that it refers directly or indirectly to one of the objects you want to store. See page 104.
- 4 Commit the transaction. This stores the object that the database root refers to and any objects that object references. See page 90.

In general, you should not create a root for each object you want to store in a database. You must create at least one root to store an object in a database by which all other objects can ultimately be reached.

How Objects Become Persistent

Objects can become persistent in several ways:

- An application assigns a transient object to a database root. ObjectStore immediately migrates the object to the default cluster of the default segment in the database. When the transaction commits, any transient objects that are reachable from the object assigned to the root are also stored in the default segment and cluster.
- A transient object is reachable from a persistent object. When the transaction commits, ObjectStore stores the reachable object in the same segment and cluster as the persistent object.
- An application invokes the `ObjectStore.migrate()` method on a transient object and specifies a particular database, segment, or cluster. When the transaction commits, ObjectStore stores the migrated object in the specified location. ObjectStore also stores in the same location any transient objects that are reachable from the migrated object.

Storing Objects in a Particular Segment or Cluster

When you want to store objects in a particular segment or cluster that is not the default segment or cluster, use the following steps:

- 1 Open a database and start a transaction.
- 2 If you have not already created the segment or cluster in which you want to store the objects, create the segment with the `Database.createSegment()` or the `Segment.createCluster()` method.
- 3 Invoke `ObjectStore.migrate()` to store an object and specify the segment or cluster in which you want to store the object.
- 4 If there are other transient objects that you want to store in the same segment or cluster, you can either migrate them as well or you can make those objects reachable from the migrated object.
- 5 Commit the transaction.

`ObjectStore` stores all transient objects that are reachable from the migrated object in the same segment and cluster as the migrated object.

To allow access to the objects, they must also be reachable from a database root.

Caution

If you are going to have objects in one segment that refer to objects in another segment, it is crucial that you plan before you store any objects. Grouping Objects in Multiple Clusters and Segments on page 77 provides information about the issues you should consider. It is important to familiarize yourself with this information because doing so can help you avoid problems.

In addition, if you have objects in one database that refer to objects in another database, make sure that the databases are affiliated. See Referencing Objects Stored in Other Databases on page 117 for more information.

What Is Reachability?

An object `B` is considered to be reachable from object `A` when `A` contains a reference to `B`, except when the reference is from a variable marked with Java's `transient` keyword. `B` is also reachable from `A` when `A` contains a reference to some object and that object contains a reference to `B`. There are no limits to levels of reachability.

Situations to Avoid

When a transaction commits, you must ensure that objects in different segments do not refer to the same transient object. If this situation exists,

ObjectStore cannot determine the segment in which to store the transient object. Consequently, ObjectStore signals `AbortException`, aborts the transaction, and informs you that an unexported object is referred to by an object in another segment.

To avoid this situation, you must explicitly migrate transient objects to a database segment when those transient objects are referred to by objects in more than one segment. You must do this before the application commits the transaction and you must export the objects when you migrate them.

When a transaction commits, ensure that any transient objects that are reachable from persistent objects are persistence capable. If one such object is not, ObjectStore signals `AbortException`, aborts the transaction, and informs you that a reachable object is not persistence capable.

Storing Java-Supplied Objects

Some Java-supplied classes are persistence capable, while others are not persistence capable and cannot be made persistence capable. A third category of classes can be made persistence capable, but there are important issues to consider when you do so. Be sure to read *Java-Supplied Persistence-Capable Classes* on page 311.

Retrieving Persistent Objects

To read the contents of an object in a database, you must first obtain a reference to an object. There are several ways you can do this:

- Use a database root.
- Iterate over the objects in the segments and clusters of the database.
- Use external references.

The following sections describe these alternatives.

This section discusses the following topics:

- Steps for Retrieving Persistent Objects
- Determining the Database That Contains an Object
- Determining Whether an Object Has Been Stored
- Locking Objects

Steps for Retrieving Persistent Objects

Follow these steps to retrieve a persistent object from a database:

- 1 Open the database.
- 2 Start a transaction. If you want to modify the object, start an update transaction.
- 3 Obtain a persistent by object using a database root, an external reference, or by iterating over the objects in a segment or cluster of the database.
- 4 Access the object just as you would access a transient object.

If you do not plan to run the postprocessor, see *Making Object Contents Accessible* on page 242.

Determining the Database That Contains an Object

You can use the `Database.of()` method to determine the database in which an object is stored. The method signature is

```
public static Database of(Object object)
```

If the specified object has been stored in a database, `ObjectStore` returns the database in which it is stored. The specified object must be a persistent primary Java object or a Java peer object. If you specify a Java peer object, it can identify a persistent or transient C++ object. If the specified object is a peer object for a transient object, the method returns the transient database.

Determining Whether an Object Has Been Stored

To determine whether an object has already been stored in a database, call the `ObjectStore.isPersistent()` method. The method signature is

```
public static boolean isPersistent(Object object)
```

If the specified object has been stored in a database, `ObjectStore` returns `true`. The specified object must not be a stale persistent object. If it is, `ObjectStore` signals `ObjectException`.

Locking Objects

You can lock an object if you want to ensure that no other session can access it. This is useful when you want to ensure that a particular operation is not interrupted. See *Locking Objects, Clusters, Segments, and Databases to Ensure Access* on page 267.

Working with Database Roots

A root is a reference to an individual object. You can get by with a single root, but you might find it convenient to have more. In general, you do not want every object in the database to be associated with a root. This is bad for performance. Each root refers to exactly one object. More than one root can refer to the same object. You cannot navigate backward from the referenced object to the database root.

This section discusses the following topics:

- Creating Database Roots
- Retrieving Root Objects
- Roots with Null Values
- Using Primitive Values as Roots
- Changing the Object Referred to by a Database Root
- Destroying a Database Root
- Destroying the Object Referred to by a Database Root
- How Many Roots Are Needed in a Database?

Creating Database Roots

When you create a database root, you give it a name and you assign an object to it. The database root refers to that object, and your application can use the root name to access that object. In other words, the object that you assign to a root is the value of that root. The database root and the object assigned to the root are two distinct objects.

You must create a database root inside a transaction. Call the `Database.createRoot()` method on the database in which you want to create the root. The method signature for this instance method on `Database` is

```
public void createRoot(String name, Object object)
```

The name you specify for the root must be unique in the database. If it is not unique, `DatabaseRootAlreadyExistsException` is signaled. The object that you specify to be referred to by the root can be either transient and persistence capable, or persistent (including null). If it is not yet persistent, `ObjectStore` immediately makes it persistent. `ObjectStore` migrates it to the default segment, which is the segment returned by

`Database.getDefaultSegment()`. Objects associated with a root are automatically exported.

Example Suppose you create the variable `db` to be a handle to a database opened for update, and an object called `anObject`, and you start an update transaction. The following line creates a database root:

```
db.createRoot("My Root Name", anObject);
```

Results In the database referred to by `db`, this creates a database root named "MyRootName" and specifies that it refers to `anObject`. `ObjectStore` immediately stores `anObject` in the database referred to by `db`. When the transaction commits, `ObjectStore` stores in the database referred to by `db` any objects that are reachable from `anObject`, if they are not already in the database. If `anObject` or any object it references refers to any transient objects that are not persistence capable and you try to commit the transaction, `ObjectStore` signals `ObjectNotPersistenceCapableException`.

Multiple roots for one object More than one root can reference the same object; an object can be associated with more than one root. For example:

```
db.createRoot("Root1", anObject);
db.createRoot("Root2", anObject);
```

Retrieving Root Objects

When you retrieve a root object, you obtain a reference to the object that is the value for the root. For example, suppose you assign an `OSVector` object, `myOSVector`, to a root named `myOSVectorRoot`. When you get the value of `myOSVectorRoot` by using the `Database.getRoot()` method, you receive a reference to the `OSVector` as follows:

```
OSVector myOSVector = (OSVector)db.getRoot("myOSVectorRoot");
```

Note that `ObjectStore` does not fetch the entire contents of the `OSVector` until you actually need it. You can obtain a reference to any object in a vector. For example, to obtain a reference to the fifth element in the vector `myOSVector`, use an assignment statement like the following:

```
Object fifth = myOSVector.elementAt(5);
```

`ObjectStore` fetches only enough contents from `myOSVector` so that it can return the reference to the fifth element. This means that `ObjectStore` has not yet fetched the contents of the elements in the vector. `ObjectStore` fetches the data for an element in a vector only when you try to access its contents.

If you develop your application by using the `ObjectStore` class file postprocessor, this delayed fetching is usually automatic. For more information on using the postprocessor, see Chapter 8, *Generating Persistence-Capable Classes Automatically*, on page 195.

In some cases, you might want to force `ObjectStore` to prefetch an entire graph of connected objects even though the application does not explicitly refer to each object in the graph. Use the `ObjectStore.deepFetch()` method to do this.

List of all roots

To obtain a list of the roots in a database, call the `getRoots()` method on the database. The signature of this method is

```
public DatabaseRootEnumeration getRoots()
```

Roots with Null Values

It is possible to create a root with a `null` value. This is useful for creating roots in preparation for assigning objects to them later. If you create a root with `null` or later set a root to `null`, the `getRoot()` method returns a `null` value, which indicates that there is no object associated with the root. It does *not* mean that the root does not exist. If the root does not exist, `ObjectStore` signals `DatabaseRootNotFoundException`.

Using Primitive Values as Roots

If you want to store a primitive value as an independent persistent object, such as the value of a root, use an instance of a wrapper class, such as an `Integer`. For example:

```
db.createRoot("foo", new Integer(5));
```

This assigns the value 5 to the root named `foo`.

You cannot directly store primitive values in a database. You can define a primitive value as a field in a persistence-capable object.

Changing the Object Referred to by a Database Root

After you create a database root, you can change the object that it refers to. Inside an update transaction, call the `Database.setRoot()` method on the database that contains the root. You must specify an existing root and you can specify either a transient (but persistence-capable) or a persistent object. If you specify a transient object, `ObjectStore` immediately stores it in the default segment of the database. The default segment is the segment

returned by the `Database.getDefaultSegment()` method. The method signature for changing the object associated with a root is

```
public void setRoot (String name, Object object)
```

If `ObjectStore` cannot find the specified root, `DatabaseRootNotFoundException` is signaled.

Destroying a Database Root

To destroy a database root, call the `destroyRoot()` method on the database that contains the root that you want to destroy. An update transaction must be in progress. Specify the name of the root. If `ObjectStore` cannot find the specified root, it signals `DatabaseRootNotFoundException`. The method signature is

```
public void destroyRoot (String name)
```

This has no effect on the referenced object except that it is no longer accessible from that root. It might still be the value of another root, or it might be pointed to by some other persistent object. If a value of a root is no longer referenced after the root is destroyed, the object becomes unreachable. You can invoke `ObjectStore.destroy()` on it while you still have a reference to it. Alternatively, you can run the persistent GC to remove all unreachable objects. See *Performing Garbage Collection in a Database* on page 61.

Destroying the Object Referred to by a Database Root

If you want to destroy the object that a database root refers to and you want to continue to use that database root, you also must set the root to refer to `null` or to another object. If you do not do this and you try to use the root, `ObjectStore` signals `ObjectNotFoundException`. This is because the root refers to a destroyed object. For example, the correct sequence is similar to the following:

```
Object object = db.getRoot("username");
db.setRoot("username", null);
ObjectStore.destroy(object);
```

How Many Roots Are Needed in a Database?

It is important to realize that you need not create a root for most objects that you want to store in a database. You need to create roots only for top-level objects that you want to look up by name. You must have at least one root to be able to navigate through a database. Without a root, you have no way of accessing the objects in the database.

Think of a database root as the root of a tree. From the root, you can climb the tree. For many applications, a root is some kind of container, such as an instance of `OSTreeMapString` or `OSVector`, or an array. After you create one or more database roots, you create other objects that are referred to by fields of the objects that the roots refer to. These objects become persistent when you commit the transaction in which you create them. In a subsequent transaction, you can look up the root objects by name and navigate from them to any other reachable persistent objects.

Too many roots can cause performance problems. The maximum practical number of roots within a database is approximately 100. Databases store roots in a vector and `ObjectStore` uses a linear search to find roots. To avoid long look-up times, restrict the number of database roots.

Iterating Through the Objects in a Cluster, Segment, or Database

To obtain an iterator for the objects in a segment, call the `Segment.getObjects()` method. To obtain an iterator for the objects in a cluster, call the `Cluster.getObjects()` method.

These methods give you access to any objects that are unreachable but that have not yet been garbage collected. The methods also provide an application-independent means of processing all objects within a segment or cluster. The method signature for both methods is

```
public Iterator getObjects()
```

This method returns an iterator object. After you have this object, you can use the following methods to iterate through the objects:

- `Iterator.next()`
- `Iterator.hasNext()`

The `Segment.getObjects()` and `Cluster.getObjects()` methods have an overloading that takes a `java.lang.Class` object as an argument and returns an iterator over all objects of that type in the database. The type can be an interface, a class, or an array type.

If your session or another session adds an object to a segment or cluster after you create an iterator, the iterator might or might not include the new object. If it is important for the iterator to include all objects accurately, you should create the iterator again.

After you create an iterator by using `Segment.getObjects()` and `Cluster.getObjects()`, objects in the segment or cluster might be destroyed. When you use an iterator to iterate through the objects, `ObjectStore` skips any destroyed objects. Note that this means that the `hasNext()` method might change the state of the iterator to skip destroyed objects. If you destroy a `com.odi.coll` dictionary from C++, `ObjectStore` does not recognize that the dictionary has been destroyed. In this case, the enumeration might return a reference to garbage.

You can use an iterator returned by the `getObjects()` method across transactions. If a transaction in which you use the iterator aborts, the iterator becomes stale and can no longer be used.

After you create an iterator for the objects in a segment or cluster, other sessions are blocked from destroying that segment or cluster until you end your transaction. If you create the iterator and then destroy the segment or cluster, the next call to `next()` or `hasNext()` causes `ObjectStore` to signal `SegmentNotFoundException` or `ClusterNotFoundException` respectively.

Using External References to Stored Objects

Typically, if you want to access a persistent object stored in an `ObjectStore` database, you must open the database and have a session and a transaction in progress, unless you committed (or aborted) your previous transaction using `ObjectStore.RETAIN_READONLY` or `ObjectStore.RETAIN_UPDATE`.

An external reference is an `ExternalReference` object that represents a reference to a persistent object stored in an `ObjectStore` database.

External references allow you to refer to a persistent object outside a transaction or a session, without opening the database. Even if you cannot access the contents of objects, it is often useful to refer to persistent objects and to pass references to those persistent objects.

For example, you must use a string representation of an external reference when your application

- Passes a reference to an object from one session to another
- Stores a reference in an ASCII file
- Transmits a reference over a serial network connection

For information on the way to create string representations for external references, see [Encoding External References as Strings](#) on page 112.

External references can be especially useful when you write a distributed application server that processes requests for many clients. This includes client/server applications that are based on Java RMI (Remote Method Invocation), ObjectStore ObjectForms, or the Object Management Group's Common Object Request Broker Architecture (CORBA).

Be careful of creating external references to objects that might be destroyed or garbage collected before the external reference is used. The garbage collector has no knowledge of an external reference to an object. So when the object being referenced is garbage collected, a tombstone is not left as a placeholder for the object. Therefore, any subsequent dereferencing of the external reference can have unpredictable results. To avoid this situation, Technical Support recommends that you export the referenced object by using `ObjectStore.export()`.

ObjectStore provides the `ExternalReference` class to represent external references. To help you use external references, this section discusses

- Creating External References
- Obtaining Objects from External References
- Encoding External References as Strings
- Using the ExternalReference Field Accessor Methods
- External References and Exported Objects
- External Reference Equality
- Reusing External Reference Objects
- External Reference Examples

Creating External References

When you create an external reference, you are creating an `ExternalReference` object that represents a reference to a persistent object stored in an ObjectStore database. You cannot create an external reference to a transient object except for an external reference to the Java `null` object.

An external reference identifies a referenced object by storing information about the referenced object's database, segment, cluster, and its location in the cluster. If the external reference is a reference to an exported object, a special export identifier is used instead of the object's location in the cluster.

You can create an external reference by using one of the three `ExternalReference` constructors or by parsing a string representation of an external reference.

Using the no-argument constructor

The no-argument constructor creates an `ExternalReference` object that refers to the `null` object. Creating an external reference using this constructor is equivalent to using the one-argument constructor and passing `null` as the argument.

The constructor signature is

```
public ExternalReference()
```

To use the no-argument constructor, it is not necessary to open a database or to have a session or transaction in progress.

Using the one-argument constructor

The one-argument constructor creates an `ExternalReference` object that identifies the object in the argument. The specified object must be persistent or `null`. If the specified object is a Java peer object, it must be a peer to a persistent C++ object.

The constructor signature is

```
public ExternalReference(Object obj)
```

If the specified object is not `null`, the database that contains the object must be open and a session and transaction must be in progress when the one-argument constructor is called.

Using the four-argument constructor

The four-argument constructor creates an `ExternalReference` object that identifies an object in a specific database, segment, cluster, or specific location in the cluster. You usually obtain the arguments for this constructor by extracting the corresponding fields from a previously created `ExternalReference` object.

The constructor signature is

```
public ExternalReference(Database db, int SegID, int ClusID,
    int loc)
```

A session must be in progress, although it is not necessary for the database to be open or for a transaction to be in progress when the four-argument constructor is called.

Creating an external reference from a string

The `ExternalReference.fromString()` method creates an external reference by parsing and then reconstructing the string it receives from the `ExternalReference.toString()` method.

The method signature is

```
public static ExternalReference fromString(String string)
```

A session must be in progress, although it is not necessary for the database containing the object to be open or for a transaction to be in progress when the `fromString()` method is called.

Obtaining Objects from External References

To obtain the object to which an external reference refers, use the `ExternalReference.getObject()` method.

The method signature is

```
public Object getObject()
```

The database containing the `ExternalReference` object must be open and there must be a session and a transaction in progress when this method is called. The session must be the same session in which the external reference was created.

Note

An external reference is valid only during the session in which it was created. To refer to a persistent object outside a session, you must encode the contents of the external reference in a format that is session neutral, such as a string, or explicitly extract the database name, segment number, cluster number, and location fields from the `ExternalReference` object.

Encoding External References as Strings

Using the `ExternalReference.toString()` and `ExternalReference.fromString()` methods, you can create an external reference that can be used outside the session in which it was created. When the methods are used together, they act as a printer and parser, respectively. They allow you to encode an external reference as a `String`, then parse the string to rebuild an equivalent external reference.

The method signature for encoding an `ExternalReference` object as a `String` is

```
public String toString()
```


The database referred to by the `ExternalReference` object need not be open and a session or a transaction need not be in progress when the `toString()` method is called.

You can pass a string encoding of an external reference to another session or process or store it in an ASCII file.

To reconstruct an `ExternalReference` object from an encoded String, call the `fromString()` method.

The method signature is

```
public static ExternalReference fromString(String ref)
```

A session must be in progress, although the database referred to by the external reference need not be open and a transaction need not be in progress when the `fromString()` method is called.

The `toString()` and `fromString()` methods are convenient to use, but they create strings that are relatively long because the strings contain the entire pathname of the database. Sometimes you can create a more compact string using accessor methods to extract and then reconstruct an equivalent external reference using the values from the fields in an `ExternalReference` object.

Using the ExternalReference Field Accessor Methods

By extracting the necessary fields from an `ExternalReference` object, you can create a more compact string representation of an external reference that you can use outside a database, process, or session.

Get-accessor methods

To extract the fields of an `ExternalReference` object, you can use the following get-accessor methods:

- `public Database getDatabase()`
- `public int getSegmentId()`
- `public int getClusterId()`
- `public int getLocation()`

You can store the values of these fields and use them later to reconstruct the external reference.

Note that the `Database` object returned by the `getDatabase()` method is valid only during the session that created the external reference. To encode the `Database` object so that you can use it outside a session, you must translate the database into another form such as a pathname by calling the

`Database.getPath()` method. This is exactly what the `ExternalReference.toString()` method does.

Large
numbers of
external
references

If you have a large number of external references for which the database is always the same (or for which the database is known), you might be able to represent the external reference more compactly than you can using the `ExternalReference.toString()` method.

For example, in the case when the database is known, you might be able to represent the external reference by storing just three integers representing the segment identifier, cluster identifier, and the location in the cluster. After you have extracted the database, segment, cluster, and location information, you can use it to reconstruct an external reference.

To reconstruct an external reference, you can use the four-argument constructor or call the set-accessor methods for the corresponding fields:

Set-Accessor
methods

- `public void setDatabase(Database d);`
- `public void setSegmentId(int segId);`
- `public void setClusterId(int clustId);`
- `public void setLocation(int loc);`

For example, you can create an external reference that refers to the `null` object using the no-argument constructor, then use the set methods to make the reconstructed external reference refer to a specific object. You can use the set methods repeatedly to update an `ExternalReference` object so that it refers to a variety of objects.

Note

The values for `SegmentId`, `ClusterId`, and `Location` are valid across session boundaries, but these values can become invalid if the referenced object is deleted or garbage collected.

External References and Exported Objects

When you create an external reference to an object that has not been exported, you must be aware of the following situations:

- Your application might garbage collect the object.
- Your application might delete the object and then garbage collect the tombstone.
- `ObjectStore` might move the object when reorganizing the segment.

If you try to use an external reference after one of these occurrences, your results are incorrect unless the object was exported before the external reference was created.

When you create an external reference for an exported object, `ObjectStore` stores the object's export identifier in the location field of the external reference. Exporting an object prevents the object or its tombstone from being garbage collected. It also allows `ObjectStore` to find the object if the object is moved when a segment is reorganized.

External Reference Equality

Two external references are considered to be equal if they both refer to the same object. In other words, if you call the `ExternalReference.getObject()` method on each external reference, both calls return identical objects.

You call the `ExternalReference.equals()` method to determine whether two external references refer to the same object.

The method signature is

```
public boolean equals(Object obj)
```

However, two external references can refer to the same object even though the `ExternalReference.equals()` method returns `false`. This happens when one external reference was created before the object was exported and the other external reference was created after the object was exported.

In this case, the two external references refer to the same object, but one external reference uses the exported form of the object's location; the other external reference does not. When this happens, you must use the `ExternalReference.getObject()` method on both external references, then compare the two objects to see if they are identical (if `o1==o2`).

You need not open the database or have a transaction in progress when `ExternalReference.equals()` is called.

Reusing External Reference Objects

After you create an `ExternalReference` object, you can reuse it any number of times. By reusing `ExternalReference` objects, you avoid the overhead of storage allocation and garbage collection when you use large numbers of external references.

You can modify an external reference so that it refers to a different object by using the set-accessor methods or the `setObject()` method.

The method signature is

```
public void setObject(Object obj)
```

External Reference Examples

In the following code fragments, an external reference to `myObj` is created three ways:

- As an external reference
- As an encoded `String`
- By extracting the field values from an external reference, then reconstructing it

First, an external reference (`ref`) is created to `myObj` using the one-argument constructor. This external reference remains valid across transaction boundaries even if the database is closed and reopened. However, the external reference can be used only during the same session in which it was created.

Next, the external reference (`ref`) is encoded as a `String` so that it can be used across session boundaries, then field values are extracted from the external reference to create a more compact representation of the external reference.

Finally, the external reference is resolved to obtain `myObj` using the `getObject()` method for the three types of external references.

```
// Creating an ExternalReference
ExternalReference ref = new ExternalReference(myObj);

// Encoding the ExternalReference as a string
String encodedStr = ref.toString();

// Extracting the fields from the ExternalReference object
Database refDb = ref.getDatabase();
String dbPath = refDb.getPath();
int segId = ref.getSegmentId();
int clustId = ref.getClusterId();
int loc = ref.getLocation();

//...other code...

// If the same session is still active, you can obtain the
// object from the original ExternalReference object:
Object obj1 = ref.getObject();
```

```
// If the session is no longer active, you can use the
// encoded string to obtain the object:
ExternalReference ref2 =
    ExternalReference.fromString(encodedStr);
Object obj2 = ref2.getObject();

// Or, you can use the field accessors methods
// to obtain the object:
ExternalReference ref3 = new ExternalReference();
ref3.setDatabase(Database.open(dbPath,
    ObjectStore.READONLY));
ref3.setSegmentId(segId);
ref3.setClusterId(clustId);
ref3.setLocation(loc);
Object obj3 = ref3.getObject();
```

Referencing Objects Stored in Other Databases

For objects in a database to reference objects stored in other databases, your application must call the `Database.affiliate()` method. This method must be called once for every database that is to be referenced.

This process known as *affiliation* adds the path of the target database to the pathname pool of the calling database. The two databases are known as *affiliated databases*.

Calling the `Database.affiliate()` method starts and commits a transaction if one is not already in progress. Also, it can result in concurrency conflict, because `Database.affiliate()` might block other applications from accessing the calling database until the affiliation call finishes.

A Java application needs to affiliate two databases if an object in one database refers to an object in another database.

The following example shows how to affiliate Database 2 (db2) with Database 1 (db1), so that Object 1 can reference Object 2.

```
import com.odi.*;

public class AffiliateDBs implements ObjectStoreConstants {
    public static void main(String[] args) {
        String dbname1 = args[0];
        String dbname2 = args[1];

        Session.create(null, null).join();
    }
}
```

```

Database db1 = Database.create(dbname1, 0664);
Database db2 = Database.create(dbname2, 0664);

Transaction.begin(UPDATE);

Object[] obj1 = new Object[1];
db1.createRoot("obj1", obj1);

Object[] obj2 = new Object[1];
db2.createRoot("obj2", obj2);

/* Affiliate db2 with db1 so that obj1 can refer to obj2. */
db1.affiliate(db2, false);
obj1[0] = obj2;
Transaction.current().commit();
}
}

```

Without the call to `Database.affiliate()`, the commit would fail with the signaling of a `DatabaseNotAffiliatedException` saying that an object in Database 1 refers to an object in Database 2 but that Database 2 is not affiliated with Database 1.

Updating Objects in the Database

To update objects in the database, start at a database root and traverse objects to locate the objects you want to modify. Make your modifications by updating fields or invoking methods on the object, just as you would operate on a transient object. Finally, save your changes by committing the transaction (which ends the transaction), check pointing the transaction (which starts a new transaction without losing locks), or by evicting the modified objects (which allows the transaction to remain active so the changes can be rolled back by aborting the transaction).

Whether you commit a transaction or evict an object, you can specify the state of objects after the operation. To specify the state that makes the most sense for your application, an understanding of the following background information is important:

- Background for Specifying Object State
- About Object Identity
- About the Object Table

Instructions for invoking `commit()`, `checkpoint()`, and `evict()` follow this background information.

Background for Specifying Object State

When a Java program accesses an object in an ObjectStore database, there are two copies of the object:

- The copy of the object in the database. This is the copy on the disk. It can be anything that is not a primitive. It can be a wrapper object.
- The copy of the object in your Java program. This is the copy that is referred to as a persistent object.

Normally, you need not be aware of the fact that there are two copies. Your application operates on the object in the Java program as if that is the only copy. This is the reason the documentation refers to this copy as a persistent object. However, the fact that there are two copies becomes apparent if a transaction aborts. In this case, the contents of the object in the database revert to the last committed copy. The effect of abort on the copy that is in your Java program depends on the retain mode you used for the abort.

About Object Identity

In a session, persistent objects maintain identity. Suppose there is an object in the database that is referred to by two different objects. You can reach the object in the database through two navigation paths. Regardless of the path you use, the resulting persistent object is the same object in the Java VM. In other words, if you have two unrelated objects (a and b), that refer to a third object (c), `a.c == b.c` is true.

In a single session, the Java VM never creates two distinct objects that both represent the same object in the database.

Sample class definitions

For example, suppose you have the following classes:

```
public class City {
    String name;
    int population;
}

public class State {
    City capital;
    String name;
    int population;
}
```

Creating objects

Suppose you also have the following code, which creates instances of these classes and stores them:

```
City boston = new City("Boston", 1000000);
State massachusetts = new State(
```

```
boston, "Massachusetts", 20000000);
OSHashtable cities = new OSHashtable();
cities.put("Boston", boston);
OSHashtable states = new OSHashtable();
states.put("Massachusetts", massachusetts);
db.createRoot("cities", cities);
db.createRoot("states", states);
```

This creates

- A `City` instance (Boston)
- A `State` instance (Massachusetts) with the Boston `City` instance as its capital
- Two instances of `OSHashtable` — one to hold `City` objects and one to hold `State` objects
- Two database roots — one to refer to each instance of `OSHashtable`

Accessing stored objects

Now you execute the following code to access the stored objects:

```
OSHashtable cities = (OSHashtable) db.getRoot("cities");
OSHashtable states = (OSHashtable) db.getRoot("states");
City boston1 = (City) cities.get("Boston");
State massachusetts = (State) states.get("Massachusetts");
City boston2 = massachusetts.capital;
if (boston1 == boston2)
    System.out.println("same");
else
    System.out.println("not the same");
```

Results

This code prints "same". This is because `boston1` and `boston2`, even though they are located through different paths in the database, are still represented by the same object in the Java VM and, therefore, they are `==`.

If you use `cities` to reach `boston1` and you modify `boston1`, you can then use `states` to access the updated version as `boston2`.

Strings and primitive wrappers

There are additional considerations for Strings and primitive wrapper classes.

String pooling causes some strings to be the same object even when you create them separately. If you call `new` multiple times to create multiple `String` objects, these separately created objects might actually refer to the same object when they are retrieved later in another transaction. See Description of `com.odi.stringPoolSize` on page 45. If you explicitly migrate the string to the database, it prevents `ObjectStore` from using string pooling.

A `String` or primitive wrapper object that you create with a single call to `new` might be represented by more than one persistent object when you access it

through different paths in subsequent transactions. This happens because a `String` or primitive wrapper object might be stored in the database without the overhead of a regular object. Usually, this does not matter for `Strings` and primitive wrapper objects because it is their value and not their identity that matters. If identity does matter, you can explicitly migrate wrapper objects into the database.

Identity across transactions

`ObjectStore` maintains the identity of referenced objects across transactions within the same session. The following code fragment, displaying "same", provides an example of this:

```
public
class Person {
    // Fields in the Person class:
    String name;
    int age;
    Person children[];
    Person father;
    // Constructor:
    public Person(String name, int age,
        Person children[], Person father) {
        this.name = name;
        this.age = age;
        this.children = children;
        this.father = father;
    }

    static void testIdentity() {
        // Omit open database calls

        Transaction tr = Transaction.begin(ObjectStore.UPDATE);
        Person children[] = { null, null };
        Person tim = new Person("Tim", 35, children, null);
        Person sophie = new Person("Sophie", 5, null, tim);
        children[0] = sophie;
        db.createRoot("Tim", tim);
        tr.commit();

        tr = Transaction.begin(ObjectStore.UPDATE);
        tim = (Person)db.getRoot("Tim");
        Person joseph = new Person("Joseph", 1, null, tim);
        tim.children[1] = joseph;
        tr.commit();

        tr = Transaction.begin(ObjectStore.READONLY);
        tim = (Person)db.getRoot("Tim");
        sophie = tim.children[0];
        joseph = tim.children[1];
        if (sophie.father == joseph.father)
            System.out.println("same");
        else
```

```
        System.out.println("not the same");  
        tr.commit();  
    }
```

About the Object Table

ObjectStore keeps a table of all objects referenced in a transaction. If you refer to the same object in the database twice (perhaps accessing the object through different paths), ObjectStore guarantees that there is only one copy of the object in your Java program. If you retrieve the same object through different paths, `==` returns `true` because ObjectStore preserves object identity.

If the system property `com.odi.disableWeakReferences` is set to `false` (the default), the references in the object table are *weak references*, which means that they do not interfere with the Java GC. If a Java program does not have any references to a persistent object (the copy in your Java program), other than through the ObjectStore object table, the object can be garbage collected. (The object in the database, of course, is not garbage collected.)

Committing Transactions to Save Modifications

When you commit a transaction, ObjectStore

- Saves and commits any modifications in the database
- Checks for transient objects that are referred to by persistent objects

If there are such objects, ObjectStore stores them in the database if they are persistence capable. This is called transitive persistence. All reachable persistence-capable objects become persistent through transitive persistence.

If the modifications contain references to objects that are not persistence capable, ObjectStore signals `AbortException`. The `AbortException.getOriginalException()` method returns the object that causes the exception.

- Sets the state of persistent objects after the transaction.

If objects were stored in the database for the first time during this transaction, the copies of these objects in your Java program are included in the group of persistent objects.

By default, persistent objects are stale after the transaction. If you do not want to make them stale, there are three ways to specify a default retain state for persistent objects after a commit operation:

- Call `Transaction.commit()` to specify a retain state that applies only to the transaction in which it is called.
- Call `Transaction.setDefaultCommitRetain()` to specify a default retain state that applies to the session in which it is called.
- Call `Transaction.setDefaultRetain()` to specify a default retain state that applies to the session in which it is called.

Method signatures

The `commit()` method has two overloadings. The first overloading takes no argument. The method signature is

```
public void commit()
```

The second overloading has an argument that specifies the state of persistent objects after the commit operation. The method signature is

```
public void commit(int retain)
```

The retain state only applies to the transaction in which the commit was called. You can specify the following retain states after a commit:

- `ObjectStore.RETAIN_HOLLOW`
- `ObjectStore.RETAIN_READONLY`
- `ObjectStore.RETAIN_STALE`
- `ObjectStore.RETAIN_TRANSIENT`
- `ObjectStore.RETAIN_UPDATE`

The values for the `retain` argument are described in the sections that follow.

Contents

The following topics in this section describe the different retain states for persistent objects after a commit operation:

- Setting a Default Commit Retain State for a Session
- Setting Persistent Objects to a Default State
- Making Persistent Objects Stale
- Making Persistent Objects Hollow
- Retaining Persistent Objects as Readable
- Retaining Persistent Objects as Writable
- Retaining Persistent Objects as Transient

- The Way Transient Fields Are Handled
- Caution About Retaining Unexported Objects

Exceptions	If you try to commit a transaction when an object in one segment refers to an unexported object in another segment, <code>ObjectStore</code> signals <code>AbortException</code> and aborts the transaction.
Incorrect access to persistent objects	If your application commits a transaction while an annotated method is executing, your program might incorrectly access additional persistent objects after the commit. For more information about this and a work around, see Troubleshooting Access to Persistent Objects on page 148.
Synchronization	If your application needs to synchronize on a persistent object, you might want to retain a reference to that object after a transaction ends by using one of the following <code>ObjectStore</code> constants: <code>RETAIN_UPDATE</code> , <code>RETAIN_HOLLOW</code> , or <code>RETAIN_READONLY</code> . When persistent objects become stale, <code>ObjectStore</code> does not maintain their transient identity. Their synchronized states are not saved persistently.

Setting a Default Commit Retain State for a Session

You can use the following two methods to set the default retain state for persistent objects after a transaction is committed:

- `Transaction.setDefaultCommitRetain()`
- `Transaction.setDefaultRetain()`

The default retain state set by these methods is in effect for the duration of the session in which they are called.

Use `setDefaultCommitRetain()` to specify a default retain state for persistent objects when `Transaction.commit()` is called without a `retain` argument. The `setDefaultCommitRetain()` method also sets the default retain state for persistent objects when `Transaction.checkpoint()` is called without a `retain` argument. The method signature is

```
public void setDefaultCommitRetain(int retain)
```

The values you can specify for `retain` are the same values you can specify when you call `commit()` with a `retain` argument. These values are described in the sections that follow.

Use `setDefaultRetain()` to specify a default commit retain state for persistent objects when `Transaction.commit()` is called without a `retain` argument. This method also sets the default retain states for persistent

objects when `Transaction.abort()` and `Transaction.checkpoint()` are called without a `retain` argument. The method signature is

```
public void setDefaultRetain(int retain)
```

Note

When you are using either method to set a default retain state, the method that was last called overrides any default retain state that was set previously. For example, if an application calls `setDefaultCommitRetain()` first and then calls `setDefaultRetain()`, the retain state for `checkpoint()` and `commit()` is specified by the second call.

Setting Persistent Objects to a Default State

To commit a transaction and to set the state of persistent objects to the state specified by `Transaction.setDefaultCommitRetain()` or by `Transaction.setDefaultRetain()`, call the `commit()` method without any arguments. For example:

```
tr.commit();
```

Making Persistent Objects Stale

When you call the `commit()` method with no argument, `ObjectStore` makes all persistent objects stale unless a default retain state has been specified previously by either the `setDefaultCommitRetain()` or `setDefaultRetain()` method. Stale persistent objects are not accessible and their contents are set to default values. `ObjectStore` reclaims the entry in the Object Table for the stale object and the object loses its persistent identity.

If your Java program still has references to stale objects, any attempt to use those references, such as by accessing a field or calling a method on the object, causes `ObjectStore` to signal `ObjectException`. Therefore, your application must discard any references to persistent objects when it calls this overloading of `commit()`.

Objects available for garbage collection

This overloading of `commit()` also discards any internal `ObjectStore` references to the copies of the objects in your Java program. When your application makes an object stale, `ObjectStore` makes any references from the stale object to other objects null. This makes the referenced objects, which can be persistent or transient, available for garbage collection if there are no other references to them from other objects.

Stale persistent objects are not available for Java garbage collection if your Java application has transient references to them.

Accessing objects again You can reaccess the same objects in the database in subsequent transactions. To do so, look up a database root and traverse to objects from there, or reference them through hollow objects. `ObjectStore` refetches the contents of the object and creates a new active persistent object. The new object has a new transient identity and the same persistent identity as the object that became stale.

For example:

```
Foo foo = myDB.getRoot("A_FOO");
ExternalReference fooRef = new ExternalReference(foo);
ObjectStore.evict(foo, ObjectStore.RETAIN_STALE);
Foo fooTwo = myDB.getRoot("A_FOO"); // refetch from database
ExternalReference fooRefTwo = new ExternalReference(fooTwo);
// At this point (foo == fooTwo) returns false,
// but (fooRef.equals(fooTwoRef)) returns true.
```

Advantage The advantage of using `commit()` with no argument when a default retain type is *not* specified, is that it wipes your database cache clean and typically makes all transient copies of persistent data available for Java garbage collection.

Disadvantage The disadvantage of using `commit()` is that any references to these objects that your Java program holds become unusable unless a default retain state was previously specified.

Alternative method Invoking `commit(ObjectStore.RETAIN_STALE)` is the same as calling `commit()` with no argument unless a default retain state was previously specified.

Making Persistent Objects Hollow

Call `commit(ObjectStore.RETAIN_HOLLOW)` to make persistent objects (the copies of the objects in your Java program) hollow. `ObjectStore` resets the contents of persistent objects to default values.

References to these objects remain valid; the application can use them in a subsequent transaction. If a hollow object is accessed in a subsequent transaction, `ObjectStore` refreshes the contents of the object in your Java program with the contents of the corresponding object in the database.

Outside transaction An application cannot access hollow objects outside a transaction. An attempt to do so causes `ObjectStore` to signal `NoTransactionInProgressException`.

Advantage The advantage of invoking `commit(ObjectStore.RETAIN_HOLLOW)` is that any references to persistent objects that the Java application holds remain

valid in subsequent transactions. This means that it is not necessary to renavigate to these objects from a database root.

Garbage collection

Sometimes an application might retain a reference to an object and prevent Java garbage collection that would otherwise occur. It is good practice to avoid retaining references to objects unnecessarily.

Scope

If you commit a transaction with `ObjectStore.RETAIN_HOLLOW`, then commit a subsequent transaction with no `retain` argument or `ObjectStore.RETAIN_STALE`, this cancels the previous `ObjectStore.RETAIN_HOLLOW` specification. No object references are available in the next transaction. This is true regardless of whether they were previously retained.

Retaining Persistent Objects as Readable

Call `commit(ObjectStore.RETAIN_READONLY)` to retain the copies of the objects in your Java program as readable persistent objects. `ObjectStore` maintains the contents of the persistent objects that the application read in the transaction just committed. The contents of these persistent objects are as they were the last time the objects were read or modified in the transaction just committed.

If any hollow objects exist when you commit the transaction, `ObjectStore` retains these objects as hollow objects that you can use during the next transaction.

After this transaction and before the next transaction, your application can read the contents of any retained objects whose contents were also retained. The actual contents of the object in the database might be different because another process modified it. Your application cannot modify these objects. An attempt to do so causes `ObjectStore` to signal `NoTransactionInProgressException`. Your application cannot access the contents of hollow retained objects. An attempt to do so causes `ObjectStore` to signal `NoTransactionInProgressException`.

Scope

If you commit a transaction with `ObjectStore.RETAIN_READONLY`, the contents of only those persistent objects whose contents were accessed in the transaction just committed are available to you after the transaction. This is because `ObjectStore` makes all retained objects hollow at the start of the next transaction. Any cached references to persistent objects remain valid. In the new transaction, `ObjectStore` fetches the contents of a persistent object when your application requires it.

Advantage	The advantage of using <code>commit (ObjectStore.RETAIN_READONLY)</code> is that the copies of the persistent objects in your Java program remain accessible after the transaction is over. In the next transaction, any cached references to persistent objects remain valid. ObjectStore copies the object's contents from the database when you access the object.
Disadvantage	<p>The disadvantage of using <code>commit (ObjectStore.RETAIN_READONLY)</code> is that it makes more work for the Java GC, because the contents of the copies of the objects in your Java program are not cleared.</p> <p>Your program might return results that are inconsistent with the current state of the database.</p> <p>ObjectStore cannot fetch any objects outside a transaction. This makes it difficult to ensure that methods can execute without signaling an exception. However, you can call <code>ObjectStore.deepFetch()</code> in the transaction to obtain the contents of all objects you might need. Of course, this increases the risk of the Java VM's running out of memory.</p>
Serialization	If you are using Java Remote Method Invocation (RMI) or serialization, you can call the <code>ObjectStore.deepFetch()</code> method followed by <code>commit (ObjectStore.RETAIN_READONLY)</code> . This allows you to perform object serialization outside a transaction.
Retaining collections not allowed	<p>Peer objects and, therefore, collection objects have no data members, so in Java, peer objects do not appear to be connected to other objects. Even if you explicitly iterate through the elements in a collection and then commit the transaction with <code>ObjectStore.RETAIN_READONLY</code>, you cannot access the collection outside a transaction. However, if the collection elements are not themselves peer objects or collections, you can manipulate them outside a transaction, but you cannot use the collection to access them. You must explicitly read them in the transaction, retain them, then access them directly or through another nonpeer object.</p> <p>Between transactions, you might try to read an object that you thought you retained, and receive <code>NoTransactionInProgressException</code>. Often, the cause of this is that you retained a reference to the object but not the contents of the object.</p>
Trouble-shooting	In a transaction, you might read the contents of an object but not the contents of objects the first object refers to. For example, during a transaction, suppose you access a vector but not any of the elements in the vector. When you commit the transaction, the contents of vector elements are not available in the transaction and they are not retained. In other words, to be able to read

the contents of an object between transactions, you must read that particular object during the previous transaction.

To be able to read objects between transactions, you might want to call `ObjectStore.deepFetch()` on an object. This method fetches the contents of the specified object, the contents of any objects that object refers to, the contents of any objects those objects refer to, and so on for all reachable objects.

Inside a transaction, `ObjectStore` fetches the contents of objects automatically as you read the objects. Outside a transaction, if a reference to an object, but not the contents of the object, was retained, `ObjectStore` signals `NoTransactionInProgressException`.

Following is another situation in which you would receive the `NoTransactionInProgressException`:

- 1 In a transaction, you read object A.
- 2 You commit the transaction with `ObjectStore.RETAIN_READONLY`.
- 3 You start a new transaction. It does not matter whether you access object A in this transaction.
- 4 You commit this transaction with `ObjectStore.RETAIN_STALE` or without a `retain` argument.
- 5 Outside a transaction, you try to access object A and you receive the `NoTransactionInProgressException`.

You might think that because you retained A after a previous transaction, its contents are still available. This is not the case. Because nothing was retained after the second transaction, the contents of A are no longer available.

MVCC alternative

Many applications that seem to be suitable for retaining persistent objects as readable can be better implemented with the Multiversion Concurrency Control (MVCC) feature. See Chapter 10, *Controlling Concurrency*, on page 259.

Retaining Persistent Objects as Writable

Call `commit(ObjectStore.RETAIN_UPDATE)` to retain the copies of the objects in your Java program as readable and writable. `ObjectStore` maintains the contents of the persistent objects as they are at the end of the transaction.

Sometimes, the contents of an object are not available in a transaction when you expect that they are available, such as when you receive a `NoTransactionInProgressException`. For more information about what

causes this exception, see Retaining Persistent Objects as Readable on page 127.

Specifying `ObjectStore.RETAIN_UPDATE` is exactly like specifying `ObjectStore.RETAIN_READONLY`, except that if your application accesses the objects after the transaction and before the next transaction, your application can modify as well as read the objects. At the beginning of the next transaction, `ObjectStore` discards any updates made to the persistent objects between transactions.

Retaining Persistent Objects as Transient

Call `commit(ObjectStore.RETAIN_TRANSIENT)` to convert all persistent objects associated with a session into transient objects at the end of a transaction. Because these transient objects are copies of persistent objects, they no longer represent persistent objects in a database. As a result, these transient objects can be read or modified outside transactions because they are no longer associated with a database.

Unlike objects retained at the end of transactions by using the `RETAIN_READONLY` or `RETAIN_UPDATE` arguments, objects retained by using the `RETAIN_TRANSIENT` argument no longer represent persistent objects in a database. Therefore, to reacquire the corresponding persistent object that the transient object (copy) once represented, the persistent object must be fetched again from the database. As a result of refetching the persistent object, you now have two Java objects in the Java VM. One Java object is the copy of the persistent object that is no longer associated with the database; the other Java object (the one just fetched) represents the persistent object in the database. If you perform an object identity test on these two transient objects (`O1 == O2`), the test returns `false` because they are not the same object.

You should be aware that after committing a transaction by using the `RETAIN_TRANSIENT` argument, `ObjectStore` does not prevent your application from accessing a hollow object that is referenced from a Java object that was retained as transient. The fields of this hollow object will contain values such as 0 and `null`.

Garbage collection

Just like other Java objects, objects that are retained at the end of a transaction by using the `RETAIN_TRANSIENT` argument are candidates for garbage collection when they are no longer referenced by other objects.

Advantage

The `RETAIN_TRANSIENT` argument is a faster alternative to `Database.close(retainAsTransient)` because the `RETAIN_TRANSIENT`

argument does not require the database to be opened and closed, which involves deallocating and reallocating all the resources associated with the database.

Using the `RETAIN_TRANSIENT` argument is a way to copy data from a database into the Java VM. Once the data is in the Java VM, you can manipulate the data without throwing database exceptions.

Disadvantage You must code your program carefully so that it does not inadvertently manipulate a copy of a persistent object instead of the actual transient object representing the persistent object stored in the database.

The Way Transient Fields Are Handled

ObjectStore does not modify the contents of any transient-only fields unless you have explicitly defined one of the following methods to modify transient-only fields:

- `IPersistent.clearContents()`
- `IPersistent.preClearContents()`
- `IPersistent.initializeContents()`
- `IPersistent.postInitializeContents()`

The `clearContents()` and `initializeContents()` methods that are generated by the postprocessor do not modify transient-only fields.

Advantage The advantage of using `commit(ObjectStore.RETAIN_UPDATE)` is that the copies of the objects in your Java program become scratch space between transactions. You can use them to determine the possible results for a particular scenario.

Disadvantage The disadvantage is that updates are discarded automatically at the beginning of the next transaction. This can make it difficult to debug applications that use this option indiscriminately.

Caution About Retaining Unexported Objects

Suppose you commit a transaction and retain persistent objects as hollow, readable, or writable. Retained objects that are not exported can become stale inside a subsequent transaction. This can happen if you run the schema evolution tool or use the schema evolution API. If you try to access a retained object that has become stale after schema evolution, ObjectStore signals `UnexportedObjectsBecameStaleException`.

When retained objects are exported, `ObjectStore` can correctly retain them after schema evolution.

Consequently, you must follow at least one of these rules:

- Do not retain unexported objects between transactions.
- Ensure that schema evolution does not operate on your database while your application is running.
- Catch the `UnexportedObjectsBecameStaleException`, abort the transaction, and retry the transaction. This is similar to handling `DeadlockException`, except that `UnexportedObjectsBecameStaleException` does not abort the transaction.
- Catch the exception, retrieve objects from roots again, and retry the operation. You need not abort the transaction.

Evicting Objects to Save Modifications

You might want to save modifications to an object or change the state of an object without committing a transaction. The `evict()` method allows you to do this.

Method signatures

The `evict()` method has two overloadings. The first overloading takes an object as an argument. The method signature is

```
public static void evict(Object object)
```

The second overloading has an additional argument that specifies the state of the evicted object after the eviction. The method signature is

```
public static void evict(Object object, int retain)
```

This section provides the following information about evicting objects:

- Description of Eviction Operation
- Setting the Evicted Object to Be Stale
- Setting the Evicted Object to Be Hollow
- Setting the Evicted Object to Be Read-Only
- Summary of Eviction Results for Various Object States
- Evicting All Persistent Objects
- Evicting Objects When There Are Cooperating Threads

- Committing Transactions After Evicting Objects
- Evicting Objects Outside a Transaction

Description of Eviction Operation

When you evict an object, `ObjectStore`

- Saves any modifications to the object in the database but does not commit the changes. If the transaction commits, any changes are committed. If the transaction aborts, the contents of the object in the database revert to the contents following the last committed transaction in which the object was modified.
- Sets the state of the evicted object after the eviction. This affects the copy of the object that is in your Java program. The default is that the evicted object is stale after the eviction. If you do not want it to be stale, you can specify another state when you invoke `evict()`.

References to other objects

When you evict an object, `ObjectStore` does not evict objects that the evicted object references.

You might evict an object that has instance variables that are transient strings. `ObjectStore` migrates such strings to the database and stores them in the same segment as the evicted object. As part of the eviction process, `ObjectStore` evicts the just-stored string with a specification of `ObjectStore.RETAIN_READONLY`. Consequently, after the eviction, the migrated string remains readable.

If you try to evict an object when that object refers to a transient object that is not persistence capable, `ObjectStore` signals `ObjectNotPersistenceCapableException` and does not perform the eviction. The exception message provides the class of the object that is causing the problem.

If you try to evict an object when the object refers to an unexported object in another segment, `ObjectStore` signals `ObjectNotExportedException` and cancels the evict operation. Also, `ObjectStore` signals `DatabaseNotAffiliatedException` if you try to evict an object when the object refers to an object in an unaffiliated database.

Caution

If your application evicts one or more objects while an annotated method is executing, your program might incorrectly access persistent objects after the eviction. For more information about this and a work around, see [Troubleshooting Access to Persistent Objects](#) on page 148.

Setting the Evicted Object to Be Stale

When you invoke `evict(Object)`, `ObjectStore` makes the evicted object stale. `ObjectStore` resets the contents of the copy of the object in your Java program to default values and makes the object inaccessible. Any references to the evicted object are stale and your application should discard them. The copy of the object in your Java program becomes available for Java garbage collection.

Advantage The advantage of using the `evict(Object)` method is that the evicted object and all objects it refers to become available for Java garbage collection (if they are not referenced by other objects in the Java program).

Disadvantage The disadvantage is that any references to the evicted object become stale. If you try to use the stale references, `ObjectStore` signals `ObjectException`.

The effect on accessibility to the copy of the object in your Java program is therefore similar to the effect of `commit()` and `commit(ObjectStore.RETAIN_STALE)`. However, if the transaction aborts, any changes to the evicted object are discarded.

Alternative method A call to `evict(Object, ObjectStore.RETAIN_STALE)` is identical to a call to `evict(Object)`.

Setting the Evicted Object to Be Hollow

When you invoke `evict(Object, ObjectStore.RETAIN_HOLLOW)`, `ObjectStore` makes the evicted object hollow. `ObjectStore` resets the contents of the copy of the object in your Java program to default values. References to the evicted object continue to be valid.

If the application accesses the evicted object in the same transaction, `ObjectStore` copies the contents of the object from the database to the copy in your Java program. If your application modified the object before evicting it, these modifications are included in the new copy in your Java program.

Advantage The reason to use the `evict(Object, ObjectStore.RETAIN_HOLLOW)` method is that the object and all objects it refers to become available for Java garbage collection (if they are not referenced by other objects in the Java program).

Sometimes an application might retain references to an object and prevent Java garbage collection that would otherwise occur. It is good practice to avoid retaining references to objects unnecessarily.

Setting the Evicted Object to Be Read-Only

When you invoke `evict(Object, ObjectStore.RETAIN_READONLY), ObjectStore`

- Retains references to the evicted object.
- Retains the contents of the evicted object.
- Saves any changes to the evicted object.
- Internally flags the object as read but not modified. This is because any changes are already saved. If the application later modifies the evicted object in the same transaction, `ObjectStore` modifies this flag accordingly.

Garbage collection

Any changes that were made before the object was evicted are saved in the database. (Of course, if the transaction aborts, the changes are rolled back.) Therefore, if the evicted object is not referenced by other objects in the Java program, it becomes available for Java garbage collection.

Additional changes

Your application can read or modify the evicted object in the same transaction. If it does, `ObjectStore` does not have to recopy the contents of the object from the database to your program. When the application commits the transaction or evicts the object again, `ObjectStore` saves in the database any new changes to the evicted object.

It might seem strange to evict an object with

`ObjectStore.RETAIN_READONLY` and yet be able to modify the object after the eviction. The specification of `READONLY` in this context means that as of this point in time, the evicted object has been read but not modified. The changes have already been saved but not committed. The contents are still available and can be read or updated.

Advantage

The advantage of using `evict(Object, ObjectStore.RETAIN_READONLY)` is that the updated object becomes available for Java garbage collection.

Summary of Eviction Results for Various Object States

The following table shows the results of an eviction according to the value specified for the `retain` argument.

<i>Results of Eviction</i>	<code>RETAIN_STALE</code>	<code>RETAIN_HOLLOW</code>	<code>RETAIN_READONLY</code>
Object state	Stale	Hollow	Active

Results of Eviction	RETAIN_STALE	RETAIN_HOLLOW	RETAIN_READONLY
References to evicted object	Stale	Remain valid	Remain valid
Candidate for Java garbage collection	Candidate	Can be candidate	Can be candidate

Evicting All Persistent Objects

You can evict all persistent objects with one call to `evictAll()`. The method signature is

```
public static void evictAll(int retain)
```

For the *retain* argument, you can specify

- `ObjectStore.RETAIN_HOLLOW`
- `ObjectStore.RETAIN_READONLY`
- `ObjectStore.RETAIN_STALE`
- `ObjectStore.RETAIN_TRANSIENT`

When you specify any of the *retain* arguments, `ObjectStore` applies it to all persistent objects that belong to the same session as the active thread. `ObjectStore` does this in the same way that it applies a *retain* argument to one object for the `evict(object, retain)` method. Note that `evict(object, retain)` does not have `RETAIN_TRANSIENT` as a value for the *retain* argument.

`ObjectStore.RETAIN_TRANSIENT` converts all persistent objects associated with a session into transient objects after evicting all objects from the transaction. Because these transient objects are copies of persistent objects, they no longer represent persistent objects in a database. You can read or modify the objects outside transactions.

Evicting Objects When There Are Cooperating Threads

Before an application evicts an object, it must ensure that no other thread requires that object to be accessible. For example, suppose you have code like the following:

```
class C {
    String x;
    String y;

    void function() {
```



```

        System.out.println(x);
        ObjectStore.evict(this);
        System.out.println(y);
    }
}

```

Before the first call to `println()`, the object is accessible. After the call to `evict()`, the `y` field is null and the second `println()` call fails. There are more complicated scenarios for this problem that involve subroutines that call `evict()` and cause problems in the calling functions. This problem can occur in a single thread. If there are multiple cooperating threads, each thread must recognize what the other thread is doing. See *Cooperating Threads* on page 31.

It is the responsibility of the application to ensure that the object being evicted is not the `this` argument of any method that is currently executing.

Committing Transactions After Evicting Objects

In a transaction, you might evict certain objects and specify their states to be hollow or active. If you then commit the transaction and cause the state of persistent objects to be stale, this overrides the hollow or active state set by the eviction. If you commit the transaction and cause the state of persistent objects to be hollow, this overrides an active state set by eviction. For example:

```

Transaction tr = Transaction.begin(ObjectStore.UPDATE);
Trail trail = (Trail) db.getRoot("greenleaf");
GuideBook guideBook = trail.getDescription();
ObjectStore.evict(guideBook, ObjectStore.RETAIN_READONLY);
tr.commit();

```

After the transaction commits, the application cannot use `guideBook`. Committing the transaction without specifying a `retain` argument makes all persistent objects stale (unless a `retain` value other than `RETAIN_STALE` was specified by `Transaction.setDefaultCommitRetain()` or `Transaction.setDefaultRetain()`). This overrides the `RETAIN_READONLY` specification when `guideBook` was evicted.

Evicting Objects Outside a Transaction

Outside a transaction, eviction of an object has meaning only if you retained objects when you committed the previous transaction. In other words, if you invoke the `commit(retain)` method and specify a value for the `retain` argument other than `RETAIN_STALE`, you can evict retained objects outside a transaction.

If you specified `commit(ObjectStore.RETAIN_STALE)`, there are no objects to evict after the transaction commits.

If you invoked `commit()` with any other retain value, you can call `evict()` or `evictAll()` with the value of the `retain` argument as `RETAIN_STALE` or `RETAIN_HOLLOW`. If you specify `RETAIN_READONLY`, `ObjectStore` does nothing.

Outside a transaction, if you make any changes to the objects you evict, `ObjectStore` discards these changes at the start of the next transaction. They are not saved in the database.

Aborting Transactions to Cancel Changes

If you modify some objects, then decide that you do not want to keep the changes, you can abort the transaction. Aborting a transaction

- Ensures that the objects in the database are as they were just before the aborted transaction started
- Sets the state of persistent objects from the transaction

Only the state of the database is rolled back. The state of transient objects is not undone automatically. Applications are responsible for undoing the state of transient objects. Any form of output that occurred before the abort cannot be undone.

Caution

If your application aborts a transaction while an annotated method is executing, your program might incorrectly access additional persistent objects after the abort operation. For more information about this and a work around, see [Troubleshooting Access to Persistent Objects](#) on page 148.

Method signatures

The `abort()` method has two overloadings. The first overloading takes no argument. The method signature is

```
public void abort()
```

The second overloading has an argument that specifies the state of persistent objects after the abort operation. The method signature is

```
public void abort(int retain)
```

The `retain` state only applies to the transaction in which the abort was called. You can specify the following retain states after a abort:

- `ObjectStore.RETAIN_HOLLOW`

- `ObjectStore.RETAIN_READONLY`
- `ObjectStore.RETAIN_STALE`
- `ObjectStore.RETAIN_UPDATE`
- `ObjectStore.RETAIN_TRANSIENT`

The values for the *retain* argument are described in the sections that follow.

Contents

This section discusses the following topics:

- Setting a Default Abort Retain State for a Session
- Setting Persistent Objects to a Default State
- Specifying a Particular State for Persistent Objects

Setting a Default Abort Retain State for a Session

You can use the following two methods to set the default retain state for persistent objects after a transaction is aborted:

- `Transaction.setDefaultAbortRetain()`
- `Transaction.setDefaultRetain()`

The default retain state set by these methods is in effect for the duration of the session in which they are called.

Call the `setDefaultAbortRetain()` method to set the default state for persistent objects after a transaction is aborted. The default retain state is in effect for the duration of the session in which it is called. The method signature is

```
public void setDefaultAbortRetain(int newRetain)
```

The values you can specify for *newRetain* are the same values you can specify when you call `abort()` with a *retain* argument. These values are described in the next section.

Another way to set the default state for persistent objects after a transaction is aborted is to call the `setDefaultRetain()` method. The method signature is

```
public void setDefaultRetain(int retain)
```

The `setDefaultRetain()` method also sets the default retain states for `Transaction.commit()` and `Transaction.checkpoint()` when these methods are called without a *retain* argument. The default retain state for persistent objects is in effect for the duration of the session in which it is called.

Note When you are using either method to set a default retain state, the method that was last called overrides any default retain state that was set previously. For example, if an application calls `setDefaultAbortRetain(int retain)` first and then calls `setDefaultRetain(int retain)`, the retain state for `abort()` is specified by the second call.

Setting Persistent Objects to a Default State

To abort a transaction and to set the state of persistent objects to the state specified by `Transaction.setDefaultAbortRetain()` or by `Transaction.setDefaultRetain()`, call the `abort()` method without any arguments. The default state is stale if a default retain state is not specified.

The method signature is

```
public void abort()
```

For example:

```
tr.abort();
```

Specifying a Particular State for Persistent Objects

To abort a transaction and to specify a particular state for persistent objects after the transaction, call the `abort(retain)` method on the transaction. The method signature is

```
public void abort(int retain)
```

The `retain` value you specify affects the retain state of the persistent objects for the transaction in which it is called.

The following example aborts a transaction and specifies that the contents of the active persistent objects should remain available to be read:

```
tr.abort(ObjectStore.RETAIN_READONLY);
```

The values you can specify for `retain` are described next.

The rules for Java garbage collection of objects retained from aborted transactions are the same as for objects retained from committed transactions. See [Committing Transactions to Save Modifications](#) on page 122.

RETAIN_STALE `ObjectStore.RETAIN_STALE` resets the contents of all persistent objects to their default values and makes them stale. This is the same as calling `abort()` when `Transaction.setDefaultAbortRetain()` or

	<p><code>Transaction.setDefaultRetain()</code> have not been called or one has been called with <code>ObjectStore.RETAIN_STALE</code> as its argument.</p>
RETAIN_ HOLLOW	<p><code>ObjectStore.RETAIN_HOLLOW</code> resets the contents of all persistent objects to their default values and makes them hollow. In the next transaction, you can use references to persistent objects from this transaction.</p>
RETAIN_ READONLY	<p><code>ObjectStore.RETAIN_READONLY</code> retains the contents of unmodified persistent objects that were read during the aborted transaction. Any objects that were modified become hollow objects, as if <code>ObjectStore.RETAIN_HOLLOW</code> had been specified. Objects whose contents were read but not modified in the aborted transaction can be read after the aborted transaction.</p> <p>If you try to modify a persistent object before the next transaction, <code>ObjectStore</code> signals <code>NoTransactionInProgressException</code>. If you modified any persistent objects during the aborted transaction, <code>ObjectStore</code> discards these modifications and makes these objects hollow as part of the abort operation.</p> <p>During the next transaction, the contents of persistent objects that were not modified during the aborted transaction are still available.</p>
RETAIN_ UPDATE	<p><code>ObjectStore.RETAIN_UPDATE</code> retains the contents of persistent objects that were read or modified during the aborted transaction. The values that are retained are the last values that the objects contained before the transaction was aborted. Even though the changes to the modified objects are undone with regard to the database, the changes are not undone in the objects in the Java VM.</p> <p>While you are between transactions, the changes that were aborted are still visible in the Java objects. At the start of the next transaction, <code>ObjectStore</code> discards the modifications and reads in the contents from the database. Objects that were read or modified in the aborted transaction can be modified between the aborted transaction and the next transaction. If you modify any persistent objects during or after the aborted transaction, <code>ObjectStore</code> discards these modifications and makes these object hollow at the start of the next transaction.</p> <p>During the next transaction, the contents of persistent objects that were not modified during or after the aborted transaction are still available.</p>
RETAIN_ TRANSIENT	<p><code>ObjectStore.RETAIN_TRANSIENT</code> converts all persistent objects associated with a session into transient objects after aborting a transaction. Because these transient objects are copies of persistent objects, they no longer</p>

represent persistent objects in a database. They can be read or modified outside transactions.

Destroying Objects in the Database

You can explicitly destroy any object that you want to be deleted from persistent storage. You can also perform persistent garbage collection, which destroys unreachable objects in the database. See *Performing Garbage Collection in a Database* on page 61. The discussion of the destroy operation covers the following topics:

- Calling `ObjectStore.destroy()`
- Destroying Objects That Refer to Other Objects
- Destroying Objects That Are Referred to by Other Objects

Calling `ObjectStore.destroy()`

To destroy an object, call `ObjectStore.destroy()`. The method signature is

```
public static void destroy(Object object)
```

The object you specify must be persistent or the call has no effect. The database that contains the object must be open for update and an update transaction must be in progress.

If the destroyed object either implements the `IPersistent` interface or is an array, you cannot access any of its fields after you destroy it.

After you invoke `ObjectStore.destroy()` on a primary Java object, `ObjectStore` leaves a tombstone. If you try to access the destroyed object, the tombstone causes `ObjectStore` to signal `ObjectNotFoundException`.

Destroying Objects That Refer to Other Objects

By default, when you destroy an object, `ObjectStore` does not destroy objects that the destroyed object references.

There is a hook method, `IPersistentHooks.preDestroyPersistent()`, that you can define. `ObjectStore` calls this method before actually destroying the specified object. This method is useful when an object has underlying structures that you want to destroy along with the object. The default implementation of this method does nothing.

You can use `preDestroyPersistent()` to propagate the destroy operation to child objects that are referenced by the one being destroyed. If you do this, be careful that the child objects themselves are not referenced by other objects in the database. If an object attempts to use a reference to an explicitly destroyed object, `ObjectStore` signals `ObjectNotFoundException`. If you are not certain whether a specific object might be referenced elsewhere, it is better to avoid explicitly destroying the object. Let the persistent GC do the job instead.

OSHashtable and OSVector

When you delete a `com.odi.util.OSHashtable` or `com.odi.util.OSVector` object, `ObjectStore` deletes the hash table or vector and its own internal data structures. `ObjectStore` does not delete the keys or elements that were inserted into the hash table or vector. Doing so might cause problems because other Java objects might refer to those objects.

However, sometimes you want to destroy the objects in a hash table or vector as well as the hash table or vector itself. Suppose you have a class in which one of the instance variables is a `com.odi.util.OSVector`. You might want to ensure that whenever an instance of this class is destroyed, the `OSVector` and its contents are also destroyed. To do this, you can define a `preDestroyPersistent()` method on your class. Define this method to iterate over the elements in the vector, destroy each one, then destroy the `com.odi.util.OSVector`.

Types not destroyed

When you call the `ObjectStore.destroy()` method on an object, it does not destroy fields in the object that are

- **String types**
- Instances of wrapper classes that have been explicitly migrated with the `ObjectStore.migrate()` method

For additional information about `ObjectStore` treatment of `String` instances, see [Description of Special Behavior of String Literals](#) on page 315. For example, if you define a class such as the one following, when you destroy an instance of this class, you should also explicitly destroy `s` and `d`.

```
class C {
    int i;
    String s;
    Double d;
}
```

Advantages of explicit destroy

You should always consider whether or not to have `preDestroyPersistent()` call `ObjectStore.destroy()` on fields that contain `String` types, instance of wrapper classes that have been explicitly

migrated, or types that you define. The advantages of explicitly destroying objects are

- `ObjectStore` replaces large objects or arrays with a 4-byte tombstone.
- Space is freed without your having to wait for the persistent garbage collector to run.
- Without expanding the size of the database, the length of time between persistent garbage collections can be longer.

Disadvantages of explicit destroy

The disadvantages of explicitly destroying such objects are

- You must write additional code.
- There is the risk of a dangling reference if you are not careful. For example, an unanticipated `ObjectException` might prevent an object from being destroyed.
- `ObjectStore` replaces a destroyed object with a tombstone that uses 4 bytes. This can cause fragmentation. The tombstone can also cause `ObjectStore` to signal `ObjectNotFoundException`. For example, suppose you unintentionally destroy an object that is referenced by another object. When you try to dereference the reference to the destroyed object, the tombstone causes `ObjectStore` to signal `ObjectNotFoundException`.

If you do not explicitly destroy an unreferenced object, `ObjectStore` destroys it when you run the persistent GC.

You need not have `preDestroyPersistent()` call `ObjectStore.destroy()` on fields that contain primitive types.

Example

For example, suppose you have a persistence-capable class called `MyVector` that has a private field called `contents`. When an instance of `MyVector` is persistent, the `contents` field is also persistent, but a user would not have access to it because it is private. If a user calls `ObjectStore.destroy()` on an instance of `MyVector`, the operation destroys the instance but not the `contents` object.

If you are the programmer implementing the `MyVector` class, you have two choices:

- Provide a `MyVector.destroy()` method to call `ObjectStore.destroy(contents)`. If you do this, you must ensure that users of `MyVector` understand that they should not call `ObjectStore.destroy()` on an instance of `MyVector` because doing so leaves garbage in the database.

- Provide a `preDestroyPersistent()` method that calls `ObjectStore.destroy(contents)`. This choice ensures that if a user calls `ObjectStore.destroy()` on an instance of `MyVector`, the operation cleans up the private `contents` field.

Following is code that shows the second alternative:

```
public class MyVector {
    private Object[] contents;

    public addElement(Object o) {
        contents[nextElement++] = o;
    }

    public void preDestroyPersistent() {
        if (contents != null)
            ObjectStore.destroy(contents);
    }
}
```

Destroying Objects That Are Referred to by Other Objects

The usual practice is to remove references to a persistent object before you destroy that persistent object. `ObjectStore` signals `ObjectNotFoundException` when you try to access a destroyed object. It is up to you to clean up any references to destroyed objects.

If an object retains a reference to a destroyed object, `ObjectStore` signals `ObjectNotFoundException` when you try to use that reference. This might occur long after the referenced object was destroyed. To clean up this situation, set the reference in the referring object to null.

String class	A call to <code>destroy</code> on a <code>String</code> object behaves differently. When you dereference a reference to such a destroyed object, <code>ObjectStore</code> does not signal <code>ObjectNotFoundException</code> . Instead, references to the destroyed object from objects modified in the same transaction as the destroy operation continue to have the value of the destroyed object. References to the destroyed object from objects not modified in the same transaction appear as null values when an object containing such a reference is fetched.
Hash tables	You should avoid having a hash table refer to a destroyed object. It is difficult to remove a reference from a hash table after you destroy the object that it refers to. This is because the search through the hash table for the referring object might cause <code>ObjectStore</code> to try to access the destroyed object. In fact, a search for another object in the hash table might cause <code>ObjectStore</code> to access the destroyed object. The result is that the hash table look-up procedure signals <code>ObjectNotFoundException</code> and the hash table becomes useless.

Consequently, you should always remove objects from hash tables before you destroy them.

Default Effects of Various Methods on Object State

The following table summarizes the default effects of various methods on the state of hollow or active persistent objects. You should never try to invoke a method on a stale object. If you do, `ObjectStore` tries to detect it and signal `ObjectException`. `ObjectStore` can signal `ObjectException` for objects that are instances of classes that implement the `IPersistent` interface.

Unless you manually annotate your classes to make them persistence capable, you do not write `ObjectStore.fetch()` or `ObjectStore.dirty()` calls in your application. The postprocessor inserts these calls automatically as needed.

The information in the following table assumes that you are not specifying a `retain` argument with any of the methods that accept a `retain` argument.

<i>Method the Application Calls</i>	<i>Result When Invoked on a Hollow or Active Object</i>
<code>ObjectStore.fetch()</code>	Active persistent object
<code>ObjectStore.dirty()</code>	Active persistent object
<code>ObjectStore.evict()</code>	Hollow persistent object
<code>ObjectStore.destroy()</code>	Stale persistent object

<i>Method the Application Calls</i>	<i>Result</i>
<code>Transaction.commit()</code>	Persistent objects become stale.
<code>Transaction.abort()</code>	Persistent objects become stale.

Transient Fields in Persistence-Capable

Classes

This section discusses

- Behavior of Transient Fields
- Preventing `fetch()` and `dirty()` Calls on Transient Fields

See also Creating Persistence-Capable Classes with Transient Fields on page 224.

Behavior of Transient Fields

In a persistence-capable class, a field designated with the `transient` keyword behaves as follows:

- A transient field is never stored in a database.
- A transient field can be initialized in a constructor just like any other field.
- When an object is materialized from a database, a transient field has the value that the constructor gives it.
- By overriding the `postInitializeContents()` method, you can synchronize a transient field for an object when its contents are refreshed from the database.
- When an object becomes hollow or stale, a transient field is not cleared.
- If you assign the value of a persistent object to a transient field, all memory of the reference is lost when the enclosing object is garbage collected.
- If you try to access a transient field outside a transaction, `ObjectStore` signals `NoTransactionInProgressException` if the containing object is hollow or `ObjectException` if the containing object is stale.
- Committing or aborting a transaction has no effect on a transient field.

Preventing `fetch()` and `dirty()` Calls on Transient Fields

When you run the postprocessor on a class that has transient fields, you might want to specify the `-noannotatefield` option for the transient fields. This option prevents access to the specified field from causing `fetch()` and `dirty()` calls on the containing object. This is useful for transient fields when you access them outside a transaction. Normally, access to a transient field causes `fetch()` or `dirty()` to be called to allow the `postInitializeContents()` and `preFlushContents()` methods to convert between persistent and transient states.

When you specify the `-noannotatetefield` option, follow it with a qualified field name.

Avoiding `finalize()` Methods

Technical Support strongly recommends that you do not define `java.lang.Object.finalize()` methods in application classes that are persistence capable. If your persistence-capable class must define a `finalize()` method, you must ensure that the `finalize()` method does not access any persistent objects. This is because the Java GC might call the `finalize()` method outside a transaction or from a thread that does not belong to the session of the object being finalized. Such a situation causes `ObjectStore` to signal `NoSessionException` and prevents execution of the `finalize()` method.

If your class defines a `finalize()` method, the class file postprocessor inserts annotations at the beginning of the `finalize()` method that change the persistent object to a transient object. This makes it safe to access fields of the finalized object. However, if the object has not been fetched, the fields are in an uninitialized state.

Troubleshooting Access to Persistent Objects

Incorrect program behavior can happen when your program does one of the following while an annotated method is executing:

- Aborts checkpoints, or commits a transaction
- Evicts one or all objects

The general result is that your program might incorrectly access additional persistent objects after the abort, commit, checkpoint, or eviction. The specific results vary according to the `retain` setting `ObjectStore` uses for the operation, as follows:

- `ObjectStore.RETAIN_STALE` should cause `ObjectStore` to signal `ObjectException` if your program tries to access a stale object. With the optimizations, your program might be able to access stale objects, which should not happen.

- `ObjectStore.RETAIN_HOLLOW`
`ObjectStore.RETAIN_READONLY`
`ObjectStore.RETAIN_TRANSIENT`
`ObjectStore.RETAIN_UPDATE`

These settings might cause your program to retrieve `null` or `0` values in place of correct values. Also, `ObjectStore` might fail to save some modifications in the database.

The class file postprocessor (`osjcfp`) uses three optimizations that can allow incorrect access to persistent objects. You can disable these optimizations by using the following `osjcfp` options:

- `-noarrayopt` disables optimization of `fetch()` and `dirty()` calls for array objects in looping constructs. This causes `osjcfp` to make the calls to `fetch()` or `dirty()` in every iteration rather than only in the first loop iteration.
- `-nothisopt` disables optimization of `fetch()` and `dirty()` calls for access to fields relative to `this` in nonstatic member methods. This causes `osjcfp` to insert a `fetch()` or `dirty()` call for each access to a field in `this`.
- `-noinitializeropt` disables optimization of `fetch()` and `dirty()` calls in constructors. Specify this option when you want the postprocessor to perform full annotation on constructors. When you specify this option, it applies to all classes that the postprocessor makes persistence capable.

Handling Unregistered Types

`ObjectStore` creates objects of type `UnregisteredType` when it must create a persistent object and it cannot find a class file for that object. The class might not be found because of a problem with the `CLASSPATH` or because the class is not available for a particular database.

If your application receives error messages that indicate unregistered types, the information here can help you determine what is happening and what to do about it. This section discusses

- How Can There Be Unregistered Types?
- Can Applications Work When There Are Types Not Registered?
- What Does `ObjectStore` Do About Unregistered Types?
- When Does `ObjectStore` Create `UnregisteredType` Objects?

- Can Your Application Run with UnregisteredType Objects?
- Troubleshooting ClassCastExceptions Caused by Unregistered Types
- Troubleshooting the Most Common Problem

How Can There Be Unregistered Types?

How can there be a type in the database with no corresponding `ClassInfo` subclass? This can happen when

- The `CLASSPATH` environment variable has been changed since the object was stored in the database and the class is no longer in the `CLASSPATH`.
- The `CLASSPATH` might include the directory or `.zip` or `.jar` file that contains the original class files, but not the directory or `.zip` or `.jar` file that contains the postprocessed class.
- The database includes an instance of a private class and there is not a corresponding `ClassInfo` subclass that describes that class. `ObjectStore` uses the reflection API to analyze persistence-capable public classes, but it is not available for private classes. Therefore, the `osjcfp` preprocessor creates a special subclass of `ClassInfo` for private classes that must be found whenever a private class is found in a database.

Can Applications Work When There Are Types Not Registered?

In some situations, it might not matter to your application that there is an object whose type is unregistered. For example, suppose you are looking up an element in a hash table. One of the elements in the hash table is of an unregistered type, but it is not the element you are looking for. Because `ObjectStore` creates an `UnregisteredType` object instead of signaling an exception, your application can keep running.

What Does ObjectStore Do About Unregistered Types?

`ObjectStore` provides the abstract class `UnregisteredType` to represent objects whose types are unregistered. When `ObjectStore` cannot find the `ClassInfo` subclass for a type that is referenced in your application, it

- Creates an `UnregisteredType` object to represent the type
- Uses the `UnregisteredType` object in place of the hollow object it would have created

You can never read or modify an `UnregisteredType` object. Because of this, it is important for you to understand

- When `ObjectStore` creates `UnregisteredType` objects
- Whether `ObjectStore` can use `UnregisteredType` objects in a particular situation

With this information, you can determine whether your application can run with objects of unregistered types. Your application can continue to run as long as you do not try to read or modify an `UnregisteredType` object.

When Does `ObjectStore` Create `UnregisteredType` Objects?

`ObjectStore` creates an `UnregisteredType` object when it encounters an object in a database and it determines that the type of that object is not registered.

`ObjectStore` encounters an object in a database when it

- Obtains the value of a database root
- Initializes an object and the value of one of the fields is a class, interface, or an array
- Initializes an array and the element type of the array is a class, an interface, or array
- Iterates over all objects in a segment

In the above list, *initialize* means to read the contents of the object out of the database and into the persistent Java object. This happens when `ObjectStore` calls `IPersistent.initializeContents()` and `IPersistent.postInitializeContents()`.

When `ObjectStore` encounters an object in a database, it determines whether there is already a Java object for the object in the database. If there is, `ObjectStore` uses that object. If there is not, `ObjectStore` checks to see whether the type of the object is registered.

If the type is not registered, `ObjectStore` tries to load the `ClassInfo` subclass for the type and register it. `ObjectStore` uses the regular Java class loading mechanism. Usually, this means that `ObjectStore` searches your `CLASSPATH`. Depending on the Java implementation you are using, Java class loading can also involve Java `ClassLoader` objects, as described in the *Java Language Specification*.

If `ObjectStore` cannot load the `ClassInfo` subclass, it cannot register the type and therefore it cannot create a hollow object for the type. In this case, `ObjectStore` creates a new Java object of type `UnregisteredType` and uses it in place of the hollow object.

Can Your Application Run with `UnregisteredType` Objects?

`ObjectStore` can use the `UnregisteredType` object if `java.lang.Object` is the type of the field in which the reference is being stored. For example, suppose you have the following class:

```
class Person {
    Pet mypet;
    Object mytrash;
}
```

You also have a database that contains one `Person` object. The value of the `Person.mypet` instance variable is an instance of the `Pet` class. The value of the `Person.mytrash` instance variable is an instance of the `Shoe` class.

Now suppose that the `Pet` class is an unregistered class. Your application opens the database and tries to read the `Person` object. This means that `ObjectStore` must initialize the `Person` object. When `ObjectStore` recognizes that the `Pet` class is unregistered, it creates an `UnregisteredType` object. `ObjectStore` then tries to assign the `mypet` instance variable to the `UnregisteredType` object. The code to do this is something such as the following:

```
mypet = (Pet) (handle.getClassField(1, XXX));
```

Typically, the postprocessor generates this code but you can specify it yourself in the `IPersistent.initializeContents()` method. In any case, the call to `handle.getClassField()` returns an `UnregisteredType` object. The cast to `Pet` is required because `Pet` is the type of the `mypet` instance variable. However, this cast does not work. You cannot cast an `UnregisteredType` to `Pet` because `UnregisteredType` is not `Pet` and is not a subclass of `Pet`. The Java VM signals a `ClassCastException` in the middle of the initialization. The `Person` object is never initialized.

Now suppose that the `Pet` class is registered and that the `Shoe` class, which is the type of the `Person.mytrash` instance variable, is not registered. `ObjectStore` creates an `UnregisteredType` object and the `handle.getClassField()` method returns it:

```
mytrash = (Object) (handle.getClassField(1, XXX));
```


This time, the cast works correctly because `UnregisteredType` is a subclass of `Object`. The initialization succeeds and the application continues to run.

Troubleshooting `ClassCastException`s Caused by Unregistered Types

If `ObjectStore` creates an `UnregisteredType` object and you do not try to do anything with it, your application should work well. Now suppose you try to do something with it. Because it exists, it must be in a variable of type `java.lang.Object`. (If it were not, you would have had trouble with it earlier, as in the `Pet` example in the previous section.)

You cannot do very much with objects of type `Object`, so it is likely that the first thing you would do is try to cast the `UnregisteredType` object to some specific type that you expect it to be. However, this does not work. If you try to cast an `UnregisteredType` object to a type other than `java.lang.Object` or `UnregisteredType`, the Java VM signals `ClassCastException`.

Unfortunately, the `ClassCastException` does not identify the type that is unregistered. There are two ways that you can determine the name of the type that is not registered:

- Change your program.
- Set the `com.odi.trapUnregisteredType` property.

Somewhere in your program, you have a variable of type `Object` whose value is an object of the `UnregisteredType` class. Modify your program to cast this variable to type `UnregisteredType`, then invoke the `getTypeName()` method on the `UnregisteredType` object. This returns the name of the type that is unregistered.

The disadvantage of this approach is that you must edit and recompile your code.

`ObjectStore` provides the `com.odi.trapUnregisteredType` property to help you determine the class that is unregistered. The default is that this property is not set, and it is usually best to use the default.

When `ObjectStore` determines that a type is not registered, it checks the setting of the `com.odi.trapUnregisteredType` property. If the property is not set (the default), `ObjectStore` creates an `UnregisteredType` object to represent the unregistered type. If `com.odi.trapUnregisteredType` is set, `ObjectStore` signals `FatalApplicationException` and provides a message indicating the name of the class that is unregistered.

The advantage of the `com.odi.trapUnregisteredType` property is that it provides the name of the class that is unregistered.

The disadvantage of the `com.odi.trapUnregisteredType` is that as soon as `ObjectStore` encounters the first object whose type is unregistered, your application stops running. If the object you want information about is the second object of an unregistered type that `ObjectStore` would encounter, `ObjectStore` never reaches that second object. When you set `com.odi.trapUnregisteredType`, `ObjectStore` signals `FatalApplicationException` as soon as it encounters the first object whose type is unregistered.

Troubleshooting the Most Common Problem

A common situation in which an `UnregisteredType` object signals `ClassCastException` occurs when you try to obtain a database root (`Database.getRoot()`) and the value of the root is an `UnregisteredType` object. For example:

```
Foo foo = (Foo) db.getRoot("foo");
```

If the `Foo` class is unregistered, the Java VM signals a `ClassCastException` when it comes to the `(Foo)` cast operation. See the previous section for two ways to determine the class that is unregistered in this situation.

However, when the value of a root is an unregistered type, it can mean that none of your persistence-capable types is registered. This is often true when an `UnregisteredType` object signals a `ClassCastException` very early in your program. Your best course of action is likely to be to ensure that your persistence-capable classes are in your `CLASSPATH` rather than trying to determine the class that is not registered.

Troubleshooting OutOfMemoryError

If you are storing many large objects, it might appear as though they are never garbage collected. In fact, you might receive the `java.lang.OutOfMemoryError`. Following is a discussion of why this might happen and what you can do.

When a transaction commits, `ObjectStore` releases references to all objects except those that are exported. `ObjectStore` uses JDK weak references to refer to exported objects. When an object is referred to only by weak references, it can be garbage collected. However, performing an explicit Java VM GC does

not necessarily cause such weakly referenced objects to be collected and their space to be freed. Typically, the Java GC frees weakly referenced objects when it needs to grow the VM heap space. So eventually you should see the storage for these objects being reclaimed.

Try running with the `-verbosegc` option to the Java VM. If the weakly referenced objects are never freed, it is likely to be for one of the following reasons:

- The application is inadvertently retaining references to these exported objects.
- The application is using a commit `retain` option other than `ObjectStore.RETAIN_STALE` or `ObjectStore.RETAIN_HOLLOW`.
- The application has disabled use of weak references when invoking the Java VM (`-Dcom.odi.disableWeakLinks=true`).

If you want the storage for a large object to be reclaimed immediately, wrap the large object in a small dummy object. Export the dummy object instead of the large object. This causes the small object to be retained as a hollow object that is referred to with a weak reference. Upon transaction commit with `ObjectStore.RETAIN_STALE` or `ObjectStore.RETAIN_HOLLOW`, the large object is available to be garbage collected almost immediately.

Chapter 7

Working with Collections

ObjectStore provides a set of persistence-capable utility collections classes in the `com.odi.util` package. These classes rely on the interfaces defined in `java.util` in the JDK 1.2 release.

ObjectStore includes another package that contains collections classes. The `com.odi.coll` package provides the API for the ObjectStore peer collections. Use these collections when you want to access C++ as well as Java. Information about these collections is in *Developing Java Applications That Access C++*.

This chapter discusses the following topics:

Description of ObjectStore Utility Collections	157
The Way to Choose a Collection	168
Using ObjectStore Utility Collections	171
Querying ObjectStore Utility Collections	173
Enhancing Query Performance with Indexes	184
Storing Objects as Keys in Persistent Hash Tables	192
Using Third-Party Collections Libraries	194

Description of ObjectStore Utility Collections

ObjectStore provides a number of utility collections interfaces and classes in the `com.odi.util` package. In addition, ObjectStore provides a query facility in the `com.odi.util.query` package.

A collection is an object that groups together other objects. It provides an effective means of storing and manipulating groups of objects and supports operations for inserting, removing, and retrieving elements.

Collections form the basis of the ObjectStore query facility, which allows you to select those elements of a collection that satisfy a specified condition. However, some collections can be queried and others cannot. Therefore, before you create a collection and store it in a database, you should consider how you plan to use a collection. When you know what you need, you can select the best persistent collection representation for your application.

To introduce you to the ObjectStore utility collections facility, this section discusses the following topics:

- Introduction to `java.util` Interfaces and Classes
- Description of `OSHashBag`
- Description of `OSHashMap`
- Description of `OSHashSet`
- Description of `OSHashtable`
- Description of `OSTreeMapxxx`
- Description of `OSTreeSet`
- Description of `OSVector`
- Description of `OSVectorList`
- Advantages of Using ObjectStore Utility Collections
- Background About Utility Collections and JDK 1.2 Collections

Introduction to `java.util` Interfaces and Classes

The `java.util.Collection` and `java.util.Map` interfaces provide methods for operating on ObjectStore collections.

- `Collection` provides methods for operating on groups of objects in which the objects might be ordered, might be duplicated, and can be queried. The internal representation of a class that implements `Collection` might be a hash table, a binary tree, or another data structure.
 - The `java.util.List` interface extends `java.util.Collection`. In collections that implement `List`, the elements are ordered and duplicates are allowed.

- The `java.util.Set` interface extends `java.util.Collection`. In collections that implement `Set`, the elements are not ordered and duplicates are not allowed.
- `Map` provides methods for operating on groups of key/value entries. Each key can map to at most one value. You cannot query collections that implement `Map`.

The `ObjectStore` utility collections facility provides the persistence-capable `java.util` classes shown in the following table. Most of these classes implement a `java.util` interface (many implement other interfaces as well).

<i>Class</i>	<i>Implements</i>
<code>OSHashBag</code>	<code>Collection</code>
<code>OSHashMap</code>	<code>Map</code>
<code>OSHashSet</code>	<code>Set</code>
<code>OSHashtable</code>	<code>None</code>
<code>OSTreeMapByteArray</code>	<code>Map</code>
<code>OSTreeMapDouble</code>	<code>Map</code>
<code>OSTreeMapFloat</code>	<code>Map</code>
<code>OSTreeMapInteger</code>	<code>Map</code>
<code>OSTreeMapLong</code>	<code>Map</code>
<code>OSTreeMapString</code>	<code>Map</code>
<code>OSTreeSet</code>	<code>Set</code>
<code>OSVector</code>	<code>Collection</code>
<code>OSVectorList</code>	<code>List</code>

Not
necessary to
postprocess
classes

You need not postprocess the classes in the utility collections facility. They are already persistence capable. If you define a subclass that extends any of these classes and you want the subclass to be persistence capable, you must either run the postprocessor on the subclass or manually annotate the subclass.Example

The `query` demo provides an example of using `ObjectStore` with utility collections. See the `README` file in the `com/odi/demo/query` directory.

Hash code

The JDK 1.2 collections interfaces specify the behavior of the `hashCode()` method on instances of the `Set`, `Map`, and `List` types. This `hashCode()` specification is based on the contents of the collection; the `hashCode` of a

collection changes depending on the elements that are added or removed. This means that it is not advisable to store an instance of a set, map, or list class in a hash table unless the set or list is immutable and will never change.

Description of OSHashBag

An `OSHashBag` is an unordered collection that allows duplicates. `OSHashBags` not only keep track of what their elements are but also of the number of occurrences of each element. As the name implies, a hash table is the internal representation for an `OSHashBag`. `OSHashBag` directly implements the `java.util.Collection` interface, so you can query instances of `OSHashBag`.

Description of OSHashMap

An `OSHashMap` is a map that allows duplicate values but not duplicate keys. Unlike `OSHashBag`, `OSHashMap` associates a key with each value in the map. When you insert a value into an `OSHashMap`, you specify the key along with the value. You can retrieve a value with a given key. The internal representation of an `OSHashMap` is a hash table. `OSHashMaps` do not allow null keys or null values.

Because `OSHashMap` implements the `Map` interface rather than the `Collection` interface, you cannot query `OSHashMaps`. However, you can query the collection views of a map: `Map.keySet()`, `Map.values()`, and `Map.entrySet()`. See [Querying Collection Views of Map Entries](#) on page 166.

The `OSHashMap.equals()` method performs value (contents) comparisons, as described by `Map.equals()`, to determine whether two `Maps` are equal. This is the only difference between `OSHashMap` and `OSHashtable`. The `OSHashtable.equals()` method compares the identities of the two objects to determine equality. The `OSHashtable.hashCode()` method generates a hash code based on object identity; it is not based on the contents of the `OSHashtable`. For information about content comparisons and identity comparisons, see [“OSHashtable and OSVector”](#) on page 167.

Description of OSHashSet

An `OSHashSet` is an unordered collection that does not allow duplicates. If you try to insert a value into an `OSHashSet` and the set already contains that value, the set remains unchanged. `OSHashSet` implements the `java.util.Set` interface. As its name implies, a hash table is the internal representation of an `OSHashSet`. Because `OSHashSet` indirectly implements `java.util.Collection`, you can query `OSHashSets`.

`OSTreeSets` are capable of storing much larger persistent collections than `OSHashSets`. However, `OSTreeSets` must be persistent; it is not possible to create a transient instance of an `OSTreeSet`. If your collection is small, an `OSHashSet` is the better choice. If your collection is large, an `OSTreeSet` performs better.

Description of `OSHashtable`

An `OSHashtable` is also an unordered collection that allows duplicates. This class has the same APIs as `java.util.Hashtable`.

`OSHashtable` associates a key with each element. When you insert an element into an `OSHashtable`, you specify the key along with the element. You can retrieve an element with a given key. While the internal representation of an `OSHashtable` is a hash table, it is a map-like structure.

Because `OSHashtable` does not implement the `java.util.Collection` interface, you cannot query `OSHashTables`. However, you can query the collection views of an `OSHashtable`. See [Querying Collection Views of Map Entries](#) on page 166.

The `OSHashtable.equals()` and `OSHashtable.hashCode()` methods perform reference (identity) comparisons and not value (contents) comparisons. This is the only difference between `OSHashtable` and `OSHashMap`. The `OSHashMap` methods perform content comparisons. For information about content comparisons and identity comparisons, see [“OSHashtable and OSVector”](#) on page 167.

By default, an `OSHashtable` allocates room for 50 elements. You can presize an `OSHashtable` to better match what your application needs. In addition, you can delay allocation of `OSHashtable` substructure, which `ObjectStore` uses to represent the `OSHashtable` until elements are actually added to the `OSHashtable`. To do this, specify the `lazy` argument to the `OSHashtable` constructor, as follows:

```
OSHashtable(int initialBufferSize, int capacityIncrement,
            boolean lazy)
```

Description of `OSTreeMapxxx`

`OSTreeMap` is based on a B-tree representation that is tuned for large persistent collections. `OSTreeMap` is an abstract class with several concrete

subclasses. In all `OSTreeMapxxx` instances, the values are objects. Each subclass uses a different type for keys, as shown in the following table:

<i>Class</i>	<i>Key Type</i>
<code>OSTreeMapByteArray</code>	<code>ByteArray</code>
<code>OSTreeMapDouble</code>	<code>Double</code>
<code>OSTreeMapFloat</code>	<code>Float</code>
<code>OSTreeMapInteger</code>	<code>Integer</code>
<code>OSTreeMapLong</code>	<code>Long</code>
<code>OSTreeMapString</code>	<code>String</code>

An `OSTreeMapxxx` is a map that allows duplicate values but not duplicate keys. Each `OSTreeMapxxx` associates a key with a value in the map. When you insert a value into an `OSTreeMapxxx`, you specify the key along with the value. You can retrieve a value with a given key. `OSTreeMapxxx`s do not allow null keys or null values.

The `OSTreeMapxxx` classes extend `OSTreeMap`, which implements `java.util.Map`. Consequently, you cannot query `OSTreeMapxxx`s. However, you can query the collection views of a map: `Map.keySet()`, `Map.values()`, and `Map.entrySet()`. See [Querying Collection Views of Map Entries](#) on page 166.

The `OSTreeMapxxx` classes are designed for large persistent aggregations. These classes allow you to iterate over the collection or query the collection without fetching any objects from the database except those that are explicitly returned to you. ObjectStore does not create hollow objects to represent the elements until they are fetched, thus reducing Java heap overhead when a subset of `OSTreeMap` is accessed. `OSTreeMap` collections can only be persistent.

Concurrency caution

You can perform read-only operations concurrently on `OSTreeMapxxx` classes from different threads in the same session, but if you perform update operations concurrently with any other operations, your results might be incorrect.

The on-line *Java API Reference* specifies for each method in the `OSTreeMap` class whether it is a read-only or an update operation.

Exported objects

Each `OSTreeMapxxx` class has a constructor for exported objects.

Description of OSTreeSet

An `OSTreeSet` is an unordered collection that does not allow duplicates. If you try to insert a value into an `OSTreeSet` and the set already contains that value, the set remains unchanged. `OSTreeSet` implements the `java.util.Set` interface. As its name implies, a balanced tree is the internal representation of an `OSTreeSet`. Because `OSTreeSet` indirectly implements `com.odi.util.IndexedCollection`, which extends `java.util.Collection`, you can query `OSTreeSets`.

The `OSTreeSet` class is designed for very large persistent aggregations. This class allows you to iterate over the collection or query the collection without fetching any objects from the database except those that are explicitly returned to you. `ObjectStore` does not even create hollow objects to represent the elements. `OSTreeSet` collections can only be persistent.

Technical Support recommends that if you are going to query a collection that contains a particularly large number of objects, you should define the collection as an `OSTreeSet` or a subclass of `OSTreeSet`. `OSTreeSet` is the only collections class for which `ObjectStore` provides the ability to add indexes. Indexes can speed up queries on very large collections. For more information, see *Enhancing Query Performance with Indexes* on page 184.

Primary index `OSTreeSet` has a constructor that allows you to create an empty `OSTreeSet` that has a primary index. The method signature is

```
public OSTreeSet(Placement place,
                 Class primaryIndexElementType,
                 String primaryIndexPath)
```

A primary index is used for queries and for looking up objects in the `OSTreeSet`. The primary index must contain no duplicate keys and must contain all elements in the `OSTreeSet`. The benefits of a primary index include

- Faster look-up times for objects in some cases
- Faster insertion and removal of objects from the set
- An iterator that returns objects in their primary key order
- Less storage space used when compared to an `OSTreeSet` with a nonprimary index.

A primary index saves storage space *only* if an existing index is designated as a primary index. If you need to create an additional index in order to designate a primary index, then no storage space is saved unless you wanted the index anyway.

Storage space is saved when an index is designated as a primary index because the hashed-based map used to locate objects is removed from the `OSTreeSet`.

Primary index maintenance If you decide to use a primary index, you must ensure that your application performs index maintenance when modifying fields of objects that affect the primary index. Otherwise, the contents of the `OSTreeSet` are not maintained and methods such as `OSTreeSet.add()`, `OSTreeSet.contains()`, and `OSTreeSet.remove()` will not work correctly. For more information on index maintenance, see *Managing Indexes and Index Values* on page 189.

Adding a primary index To designate that an existing `OSTreeSet` index be used as a primary index, call the following method:

```
public void setPrimaryIndex(Class elementType, String path)
```

ObjectStore signals `IndexException` if the specified index is not found or allows duplicate key values, or if the `OSTreeSet` contains elements that are not instances of the `elementType`.

Obtaining a primary index To obtain a primary index from an `OSTreeSet`, call the following method:

```
public IndexMap getPrimaryIndex()
```

The primary index is returned if one exists; otherwise, `null` is returned.

Specifying no primary index To specify that none of the existing indexes of an `OSTreeSet` be used as a primary index, call the following method:

```
public void noPrimaryIndex()
```

When you specify no primary index, you are not removing the index itself. The index is still available for queries. It just means that OSJI will use a map of object hash codes to find objects in the set instead of using the primary index. This method does nothing if there is no primary index on the set.

Comparing `OSTreeSet` and `OSHashSet` The primary difference between `OSTreeSet` and `OSHashSet` is the internal representation. Also `OSTreeSet` supports indexes, whereas `OSHashSet` does not. `OSTreeSets` can only be persistently allocated. It is not possible to create a transient `OSTreeSet`. For more information on which collection to use, see *The Way to Choose a Collection* on page 168

Concurrency caution You can perform read-only operations concurrently on the `OSTreeSet` class from different threads in the same session, but if you perform update operations concurrently with any other operations, your results might be incorrect. Performing queries on `OSTreeSet` classes with methods from the `com.odi.util.query.Query` class are considered read-only operations.

The on-line *Java API Reference* specifies for each method in the `OSTreeSet` class whether it is a read-only or an update operation.

Exported objects

The `OSTreeSet` class has a constructor for creating exported objects.

The `OSTreeSet` class has a constructor for creating exported objects.

Description of `OSVector`

An `OSVector` is a persistent expandable array that implements `java.util.Collection`. You can query `OSVectors`.

An `OSVector` associates each element with a numerical position based on insertion order. By default, `OSVectors` allow duplicates. In addition to simple insert (insert into the beginning or end of the collection) and simple remove (remove the first occurrence of a specified element), you can insert, remove, and retrieve elements based on a specified numerical position or based on a specified iterator position. An `OSVector` does not have quick look-up by object or key. Therefore, the overhead for an `OSVector` is lower than for utility collections that have quick look-up.

The `OSVector.equals()` and `OSVector.hashCode()` methods perform reference (identity) comparisons and not value (contents) comparisons. This is one difference between `OSVector` and `OSVectorList`. The `OSVectorList` methods perform content comparisons. For information about content comparisons and identity comparisons, see “`OSHashtable` and `OSVector`” on page 167.

By default, an `OSVector` allocates room for 32 elements. You can presize an `OSVector` to better match what your application needs. In addition, you can delay allocation of `OSVector` substructure, which `ObjectStore` uses to represent the `OSVector` until elements are actually added to the `OSVector`. To do this, specify the `lazy` argument to the `OSVector` constructor, as follows:

```
OSVector(int initialBufferSize, int capacityIncrement,
        boolean lazy)
```

Description of `OSVectorList`

An `OSVectorList` is a collection that implements a persistent expandable array. It implements the `java.util.List` interface and functions exactly like an `OSVector`, except in the following way.

The `OSVectorList.equals()` and `OSVectorList.hashCode()` methods perform value (contents) comparisons and not reference (identity)

comparisons. This makes `OSVectorList` unsuitable for storage in a persistent hash table or any other hash-table-based collection representation. The `OSVector` methods perform identity comparisons. For information about content comparisons and identity comparisons, see “OSHashtable and OSVector” on page 167.

Advantages of Using ObjectStore Utility Collections

The advantages of using `com.odi.util` interfaces and classes are as follows:

- The interfaces and classes in `com.odi.util` rely on and extend the interfaces defined in the JDK 1.2 release.
- The classes are persistence capable.
- There are collection representations that support queries.
- Some of the classes — `OSTreeMapxxx` and `OSTreeSet` — support very large aggregations.

Querying Collection Views of Map Entries

The `OSHashMap` and `OSTreeMapxxx` classes extend `java.util.Map` and not `java.util.Collection` and, therefore, you cannot use the ObjectStore query facility on them. However, each of the classes that implements `Map` defines the following methods:

- `keySet()` returns a `java.util.Set` view of the keys contained in the map.
- `values()` returns a `java.util.Collection` view of the values contained in the map.
- `entries()` returns a `java.util.Set` view of the key/value mappings contained in the map.

The `OSHashtable` class, although it does not implement `Map`, also defines these methods.

You can use the ObjectStore query facility to query the `Collection` and `Set` views returned by the `keySet()`, `values()`, and `entries()` methods.

Transient views

While `OSHashtable`, `OSHashMap`, and the `OSTreeMapxxx` subclasses are persistence capable, the views returned by the `entries()`, `keySet()`, and `values()` methods are not. These are transient views of persistence-capable classes.

Background About Utility Collections and JDK 1.2

Collections

Following is background information about how the ObjectStore utility collections fit with the JDK 1.2 collections. This discussion assumes that you are familiar with the JDK 1.2 collections API.

ObjectStore provides a collections package that relies on and extends the JDK 1.2 `java.util` collections. In addition, ObjectStore includes querying and indexing facilities. The new collections implementations are in the `com.odi.util` package.

The core collections interfaces defined in the JDK 1.2 `java.util` package are

- `java.util.Collection`
- `java.util.Set`
- `java.util.List`
- `java.util.Map`

In the JDK 1.2, collections classes and behaviors are based on these interfaces. Consequently, you can usually use any representation that is parallel to a particular interface. The `java.util` implementations and their corresponding ObjectStore implementations are shown in the following table:

<i>Interface</i>	<i>java.util Class</i>	<i>ObjectStore Class</i>
Collection	None	<code>com.odi.util.OSTHashBag</code>
Set	<code>java.util.HashSet</code>	<code>com.odi.util.OSTHashSet</code>
Set	<code>java.util.ArraySet</code>	<code>com.odi.util.OSTTreeSet</code>
List	<code>java.util.Vector</code>	<code>com.odi.util.OSVector</code>
List	<code>java.util.ArrayList</code>	<code>com.odi.util.OSVectorList</code>
List	<code>java.util.LinkedList</code>	None
Map	<code>java.util.Hashtable</code>	<code>com.odi.util.OSTHashtable</code>
Map	<code>java.util.HashMap</code>	<code>com.odi.util.OSTHashMap</code>
Map	<code>java.util.ArrayMap</code>	None
Map	<code>java.util.TreeMap</code>	<code>com.odi.util.OSTreeMapxxx</code>

OSTHashtable
and
OSVector

`com.odi.util.OSTHashtable` and `com.odi.util.OSVector` have been updated to be parallel to most of the JDK 1.2 specifications. They do not quite meet the description of the JDK 1.2 behavior for `equals()` and `hashCode()`.

The JDK 1.2 changed this behavior in an incompatible way for these two classes.

The JDK 1.2 `List`, `Set`, and `Map` interfaces mandate an `equals()` method that does value comparison and not reference comparison. That is, two `Sets` are equal if they have the same elements, two `Lists` are equal if they have the same elements in the same order, and two `Maps` are equal if they have the same key/value pairs.

This places corresponding constraints on the `hashCode()` method because `(a.equals(b)) => (a.hashCode()==b.hashCode())`. The `ObjectStore` `OSHashtable` and `OSVector` classes, however, implement persistent (unchanging) `hashCodes` and rely on `Object.equals()`. The JDK definition for `hashCode` means that classes that meet the JDK 1.2 specification should not be stored in hash tables because their `hashCodes` change when elements are added or removed. For these two classes, `ObjectStore` retains the old identity-based definitions rather than moving to the new content-based definitions of `equals()` and `hashCode()`.

Collection interface

There are no concrete implementations of the `Collection` interface in the JDK 1.2. `Collection` is essentially a `Bag`, that is, a `Set` that might contain duplicates. `ObjectStore` includes the `com.odi.util.OSHashBag` and `com.odi.util.OSVector` classes to implement `Collection`.

The Way to Choose a Collection

Your choice of how to implement a collection depends on

- The amount of data to be stored in the collection

The following numbers can help you determine the type of collection to use. The efficiency of a collection is based on its performance and storage size.

- Java arrays are efficient for up to 1000 elements.
- `OSHashtables` are efficient for small collections that have fewer than 1000 elements.
- `OSVectors` are efficient for medium collections that have as many as 1,000,000 elements.
- `OSTreeMaps` and `OSTreeSets` are efficient for medium and large collections that have between 200 and several million elements.

The recommended size for `OSVector` overlaps with the sizes for other collection types. When sizes overlap, you should use `OSVector` when the keys are contiguous integers.

The recommended size for `OSHashtable` overlaps with the size for `OSTreeMap`. When the sizes overlap, use an `OSHashtable` only if the same element is accessed multiple times within a transaction.

- Whether the queries you use can benefit from using an index

Indexes are useful when you are querying a collection that is larger than a few hundred elements and when you know in advance what fields will be queried.

`OSTreeSet` is the only collection that can have indexes, one of which you can designate as a primary index. If an `OSTreeSet` has a primary index, that index is used for determining set membership. With a primary index, inserting and removing elements is faster when compared to a set with a nonprimary index, less storage space is used, and the iterator can return elements in their primary key order.

- Your familiarity with a third-party library

You might want to use a particular library because you already know how to use it.

- Features required by your application

If you want to access your collection from C++, use the collections in the `com.odi.coll` package. For more information, see *Developing Java Applications That Access C++*.

- Importance of compatibility with the JDK 1.2 collection interfaces

The `OSHashtable` and the `OSVector` classes are not compatible with the JDK 1.2 `Map` interface because its `.equals` operator is identity based and its hash code is not based on contents. All other collections in `com.odi.util` are compatible with the JDK 1.2 API.

- Importance of compatibility between OSJI and PSE Pro

PSE Pro does not have the `com.odi.coll` collections.

For applications that already use the `com.odi.coll` collections, the `com.odi.util.OSTreeMapxxx` classes are comparable in performance to the `com.odi.coll.Dictionary_xxx` classes, while the `com.odi.util.OSTreeSet` class is comparable in performance to the `com.odi.coll.Set` class.

Comparing Collection Classes

To help you choose the right persistent collection representation for your application, the following table compares the behavior of the different collection classes in `com.odi.util`.

Table note

In the following table

- The key in an `OSVector` and `OSVectorList` is the offset.
- Elements in an `OSTreeSet` are in key value order *only* when there is a primary index.

<i>Class</i>	<i>Elements in Key Value Order?</i>	<i>Duplicate Keys Allowed?</i>	<i>Duplicate Values Allowed?</i>	<i>.contains (Object) Slow?</i>	<i>.equals() Compares By</i>	<i>Queries Allowed?</i>	<i>Size</i>
<code>OSHashBag</code>	No	Yes	Yes	Yes	Identity	Yes	Small
<code>OSHashMap</code>	No	No	Yes	Yes	Content	No	Small
<code>OSHashSet</code>	No	N/A	No	No	Content	Yes	Small
<code>OSHashtable</code>	No	No	Yes	Yes	Identity	No	Small
<code>OSTreeMapxxx</code>	Yes	No	Yes	Yes	Content	No	Medium /Large
<code>OSTreeSet</code>	No/ Yes	N/A	No	No	Content	Yes	Medium /Large
<code>OSVector</code>	Yes	No	Yes	Yes	Identity	Yes	Medium
<code>OSVectorList</code>	Yes	No	Yes	Yes	Content	Yes	Medium

Performance-Based Recommendations for Collections

This section lists additional performance recommendations that you might want to consider when choosing which collection type to use.

OSHashtable The `OSHashtable` becomes *very* inefficient when the collection grows larger than a few thousand elements. The `OSHashtable` is faster than an `OSTreeSet` *only* when the same set of elements is accessed many times within the same transaction.

OSTreeMapxxx and OSTreeSet The `OSTreeMapxxx` and `OSTreeSet` implementations use an approach that requires the fewest object materializations for accessing their elements. As a result, `OSTreeMapxxx` and `OSTreeSet` implementations have the fastest initial access to elements, but they are not as fast as the other types of

collection classes when you are accessing the elements for the second time in a transaction.

OSTreeSet	Designating one of the indexes of an <code>OSTreeSet</code> to be a primary index makes inserting and removing elements faster.
<code>com.odi.coll</code> package	The collections in the <code>com.odi.coll</code> package are approximately the same as the <code>com.odi.util</code> collections in their look-up speed for an object, but the <code>com.odi.coll</code> collections are much slower when you are inserting and iterating through elements. You should use the <code>com.odi.coll</code> collections <i>only</i> if your Java applications will access C++ code.

Using ObjectStore Utility Collections

To help you use ObjectStore utility collections, this section discusses the following topics:

- Creating Collections
- Navigating Collections with Iterators
- Performing Collection Updates During Iteration

Creating Collections

Each collection representation has one or more constructors that you can use to create collections. For details about each class's constructors, see the *Java API Reference*. For example:

```
Database db = Database.create(args[1], ALL_READ | ALL_WRITE);
Transaction.begin(UPDATE);
db.createRoot("collection", new OSTreeSet(db));
Transaction.current().commit();
```

Navigating Collections with Iterators

The `Iterator` and `ListIterator` interfaces help you navigate within a utility collection. An *iterator* is an instance of the `java.util.Iterator` or `java.util.ListIterator` interface. It designates a position in a collection. You can use iterators to traverse collections as well as to remove elements from collections.

With the JDK 1.2, `Iterator` takes the place of `Enumeration`. `Iterator` provides the same capabilities as `Enumeration` (though method names are

different), and it allows you to remove elements from the underlying collection.

The `ListIterator` interface extends the `Iterator` interface. A class that supports traversal by `ListIterator` must also implement `List`. The additional methods that `ListIterator` provides allow you to

- Insert objects relative to the current position of the iterator
- Traverse the list in reverse as well as forward
- Replace an element in the underlying list
- Retrieve the index of an element

The `IndexIterator` interface in `com.odi.util` extends `java.util.Iterator` and allows you to traverse an index or map structure. You can use the `IndexIterator` interface to obtain the key and value for elements in the underlying collection.

Performing Collection Updates During Iteration

While you are iterating through a collection, you can use the `Iterator` and `ListIterator` interface methods to modify that collection. This assumes that the implementation of the `Iterator` or `ListIterator` interface supports the methods that modify underlying collections. (The JDK 1.2 defines some of these methods as optional. You should check the API reference information for the particular class you are using to determine exactly the behaviors that are supported.)

When a thread is iterating over a collection, that thread and cooperating threads can modify the object returned by the iteration. If you are using an `Iterator`, your application cannot add elements to the collection or change the order of the collection. If you are using a `ListIterator`, your application can only use `ListIterator` methods to modify the collection.

Suppose you do add an element in the middle of an iteration, and then try to use the same iterator. ObjectStore recognizes that the collection has been modified and signals `ConcurrentModificationException`. At this point, if you create a new iterator, it recognizes the updated collection and does not signal an exception.

Querying ObjectStore Utility Collections

The `com.odi.util.query.Query` class provides a mechanism for querying collections objects that implement the `java.util.Collection` interface. A query applies a predicate expression (an expression that evaluates to a `boolean` result) to all elements in a collection. The query returns a subset collection of all elements for which the expression is true. You can query the following classes that implement the `Collection` interface:

- `OSHashBag`
- `OSHashSet`
- `OSTreeSet`
- `OSVector`
- `OSVectorList`

To accelerate the processing of queries on particularly large collections, you can build indexes on the collection. For information about indexes, see the next section, *Enhancing Query Performance with Indexes* on page 184.

This section provides the following information about queries on ObjectStore utility collections:

- Creating Queries
- Description of Query Syntax
- Sample Program That Uses Queries
- Matching Patterns in Query Strings
- Using Free Variables in Queries
- Executing Queries
- Limitations on Queries

Creating Queries

To create a query, run the `com.odi.util.query.Query` constructor and pass in a `Class` object and a query string. Following is the constructor:

```
public Query(Class elementType, String queryExpression)
```

There is also a constructor that allows you to specify a `FreeVariables` map.

The `elementType` class or interface provides the context in which the query facility interprets `queryExpression`. This must be a publicly accessible class or interface. When your application calls the `Query.select()` or

`Query.pick()` method to execute the query against a particular collection, every element of that collection must be an instance of (in the sense of `instanceof`) the `elementType` that was specified when the query was created. Any element of the collection that is not an instance of `elementType` is not returned in the query result, even if it evaluates to `true` for the predicate.

The `queryExpression` is a predicate (that is, an expression with a boolean result) that the query facility evaluates on each element of the collection. The `queryExpression` operands can be literals and names.

Literals can be of any of the Java primitive types, including the special values `true`, `false`, and `null`. Because the query expression is a `String`, you must enclose any embedded strings in escaped quotation marks, like `\\"this\\"`.

Names can consist of a single identifier or they can consist of a sequence of identifiers separated by periods. Names can be either free variables or member names (field or method names). You must explicitly specify free variables in the `freeVariables` argument of the three-argument `Query` constructor. Any name that is not a free variable is interpreted as a member name.

Member accesses are interpreted as accessing public members, including static members of an object of class/interface `elementType`, if possible. This interpretation works as though there were an implicit `this` argument of `elementType` at the root of the name expression. Any member access that cannot be interpreted as a member access on `elementType` is interpreted as a static access. Static accesses are resolved as if the package containing `elementType` were imported.

Queries can contain methods that take arguments. The arguments can be literals or bound variables.

Example

For example, to define a simple query:

```
Query q = new Query(Employee.class, "salary < 50000");
```

The query expression can refer to classes without specifying a package name. ObjectStore treats the query expression as if it were defined in a file in another package that has imported the package of the `Class` object that was passed to the `Query` constructor. This default package matters only for class names, though, not for member access. Only public classes and members are accessible within the query.

An application can run the example query on a specific collection with a call to the `Query.select()` method that specifies the collection to be queried as the argument. For example:

```
Query q = new Query(Employee.class, "salary < 50000");
Collection employees = (Collection)db.getRoot("employees");
Set result = q.select(employees);
```

When you create a query, you do not bind it to a particular collection. You can create a query, run it once, and throw it away. Alternatively, you can reuse a query multiple times against the same collection, perhaps with different bindings for free variables, or against different collections.

If the syntax of your query is wrong, `QueryException` is thrown at the point at which you create the query. You need not wait for the application to optimize or to execute the query. However, the query facility cannot detect incorrect free variable bindings until you specify them when you execute the query on a collection.

Sample program

The following sample program uses a query that takes a method as an argument. To run this program, you need to

1 Compile it.

```
javac QueryMethodWithArgs.java
```

2 Use the postprocessor to make the class persistence capable.

```
osjcfp -dest . -inplace QueryMethodWithArgs.class
```

3 Run it as an application with no arguments.

```
java QueryMethodWithArgs

import com.odi.*;
import com.odi.util.*;
import com.odi.util.query.*;

/**
 * An example of performing queries that use methods as
 * arguments. Instances of this class have first name and last
 * name fields. The getName(useFirst) method returns the first
 * or last name, depending on the value of the useFirst
 * argument.
 * You can use an index on the getName() method with an
 * appropriate value when evaluating the query, because its
 * argument is a constant value. */

public class QueryMethodWithArgs implements
ObjectStoreConstants {

    /* Fields */

    String first;
```

```

String last;

/**
 * Constructor.
 */

QueryMethodWithArgs(String first, String last) {
    this.first = first;
    this.last = last;
}

/**
 * Returns the first name if useFirst is true, otherwise the
    last name.
 */

public String getName(boolean useFirst) {
    return useFirst ? first :last;
}

/**
 * Include the field values in the print string.
 */

public String toString() {
    return "QueryMethodWithArgs{ first = " + first + ", last =
        " + last + "
    }" ;
}

/**
 * Main routine to run application.
 */

public static void main(String[] args) {
    Session session = Session.create(null, null);
    try {
        session.join();

        Database db = Database.create("foo.odb", 0664);
        Transaction.begin(UPDATE);
        /* Create an OSTreeSet. */
        OSTreeSet set = new OSTreeSet(db);
        db.createRoot("set", set);
        /* Add two objects to it. */
        set.add(new QueryMethodWithArgs("John", "Doe"));
        set.add(new QueryMethodWithArgs("Jane", "Doe"));
        /* Add an index on getName(false), which is an index on the
            last names of the objects. */
        set.addIndex(QueryMethodWithArgs.class, "getName(false)");
        Transaction.current().commit();
    }
}

```



```

Transaction.begin(READONLY);

set = (OSTreeSet)db.getRoot("set");

/* Create a query to look for last names equal to "Doe". */
Query query = new Query(QueryMethodWithArgs.class,
    "getName(false) == \"Doe\"");
/* Perform the query. */

Iterator iterator = query.iterator(set);

/* Print the matches. */

while (iterator.hasNext())
    System.out.println(iterator.next());

    } finally {
        session.terminate();
    }
}
}

```

Description of Query Syntax

ObjectStore performs syntax analysis of the query expression in the context of the `elementType` class or interface that is passed to the query constructor. This must be a publicly accessible class or interface, or a derived type.

When the query is executed against a particular collection using the `select()` or `pick()` method, every element of that collection must be an instance (in the sense of `instanceof`) of the `elementType` that was specified when the query was created.

The `queryExpression` is a predicate. The query is executed on a collection by evaluating this query expression on each element of the collection. However, it might not be necessary to explicitly fetch and examine all elements of the collection. This depends on the available indexes and query optimization strategy.

Supported operations

Queries on utility collections can include most Java operations, as follows:

- Arithmetic: `+` `/` `-` `*` `%`
- Bitwise: `^` `|` `&`
- Unary numeric: `~` `-`
- Unary logic: `!`
- Relational: `>` `<` `<=` `>=` `instanceof`
- Equality: `==` `!=`

- String concatenation: +
- Conditional AND, OR: && ||
- Shift operations: << >> >>>
- Cast operations: (type)

Unsupported operations

The following operations are not supported:

- Assignment: = += *= /= %= -= <<= >>= >>>= &= ^= |=
- Conditional: ?:
- Array dereference: []
- New: new
- Prefix/Postfix: ++ --

Statements are not permitted. Only expressions are permitted.

For details on operations and the operands, see the *Java Language Specification*.

The operators have their usual Java meanings except for the relational and equality operators when used with `String` operands. In a query expression, ObjectStore uses these operators to compare the contents of the two strings rather than their identities. Null `Strings` are considered to be less than all other values.

String literals

In a query expression, you must enclose `String` literals in escaped quotation marks. For example:

```
new Query(Foo.class, "name == \"Davis\"")
```

You can specify wildcards in query strings. You can search for substrings and perform case-insensitive searches. See Matching Patterns in Query Strings on page 179.

Wrapper objects

The query facility treats wrapper objects just as it does other `Objects`. For example, suppose you have the query expression "`A==B`". `A` and `B` refer to `Integer` wrappers. This results in an identity check on the objects. The query facility determines whether `A` and `B` both refer to the same wrapper instance. The query facility does not check that the values of `A` and `B` are equal. You can specify "`A.intValue()==B.intValue()`" to compare contents.

This behavior might change in a future release so that the query facility treats wrapper objects in the way that it treats primitives. Consequently, you should not rely on the identity check for wrapper objects.

- Other rules
- You can use parentheses to group expressions.

The precedence and associativity of the operators is the same as that for the Java language.

The entire query expression must resolve to a Boolean value.

Sample Program That Uses Queries

In the `com/odi/demo/query` directory, there is a sample program that uses `ObjectStore` utility queries. See the `README.htm` file in that directory.

Matching Patterns in Query Strings

- Specifying a pattern-matching query
- To specify a string pattern to be matched in a query, the pattern matching operator (`~~`) is used. This operator, which has greater precedence than the multiplication operator (`*`), has two arguments. These arguments must be either `Strings` or `null`. The left-hand argument specifies the text to be checked for a match. The right-hand argument specifies the pattern to be matched.
- Pattern-matching characters
- The following characters have special meanings when used in the right-hand argument of the pattern matching operator. All other characters match themselves.

Operator	Function
<code>?</code>	Matches any single character
<code>*</code>	Matches 0 or more of any character
<code>&</code>	Escape character
<code>[</code>	Reserved
<code>]</code>	Reserved
<code>(</code>	Reserved
<code>)</code>	Reserved
<code> </code>	Reserved

- Note
- The reserved characters are invalid if they are not preceded by an ampersand (`&`).

The following table shows special two-character sequences, known as escape sequences, that start with an ampersand (`&`). These escape sequences are

used to include characters literally in the pattern without their special meaning and to enable case-insensitive matching.

Note that the ampersand (&) must appear in front of every sequence. An ampersand followed by any other character is invalid. Case sensitivity in

Escape Sequence	Function
&?	Matches a question mark
&*	Matches an asterisk
&[Matches left square bracket
&]	Matches right square bracket
&(Matches left parenthesis
&)	Matches right parenthesis
&	Matches a vertical bar
&&	Matches an ampersand
&i	Enables case-insensitive matching

matching

By default, pattern matches are case sensitive. The &i escape sequence enables case-insensitive matching for an entire pattern. This escape sequence can be specified only at the start of a pattern.

Optimizing
pattern
matching

The pattern-matching operator takes advantage of any ordered indexes available on the text being matched. If the pattern starts with a character other than an asterisk (*) or a question mark (?), the query searches only the portion of the index that matches the initial constant prefix. Therefore, patterns that specify a constant prefix produce much more efficient queries.

Pattern-
matching
examples

The following pattern-matching examples use the following class:

```
public class Person {public String name;}

• Matching a name beginning with the characters Tom:
  new Query(Person.class,"name ~~ \"Tom*\");

• Matching a name ending with the characters man or burn:
  new Query(Person.class,
    "name ~~ \"*man\" || name ~~ \"*burn\");

• Matching a name using a single wildcard character with a bound variable:
  FreeVariables vars = new FreeVariables();
```

```
vars.put("var", String.class);
Query query = new Query(Person.class, "name ~~ var", vars);
FreeVariableBindings bindings = new FreeVariableBindings();
bindings.put("var", "*Gr?y");
query.select(coll, bindings);
```

- Matching a name using a case-insensitive match for `?foo`:

```
new Query(Person.class, "name ~~ \"&i&?foo\"");
```

- Matching a name using a case-insensitive match for `*foo` appearing anywhere:

```
new Query(Person.class, "name ~~ \"&i&*foo*\"");
```

- Matching a name `foo` appearing anywhere followed by `&bar`:

```
new Query(Person.class, "name ~~ \"*foo*&&bar*\"");
```

- Matching the name `(a)`:

```
new Query(Person.class, "name ~~ \"&(a&)\"");
```

Using Free Variables in Queries

Free variables are lexically the same as identifiers in the Java language. If you use free variables in your query, you must specify them in an optional third argument to the `Query` constructor. Use the

`com.odi.util.query.FreeVariables` class. This class implements the `Map` interface. In addition, it provides type checking to ensure that the keys and values are `Strings` and `Classes`, respectively. For example:

```
FreeVariables vars = new FreeVariables();
vars.put("INPUT_SALARY", Integer.TYPE);
Query q = new Query(Person.class,
    "salary>=INPUT_SALARY", vars);
```

When you execute a query, you must bind any free variables to particular values. Do this by passing an additional argument to the `Query.select()` or `Query.pick()` method. This argument must be of type

`com.odi.util.query.FreeVariableBindings`. This class, like `FreeVariables`, implements the `Map` interface and provides additional type checking to ensure that the keys are `Strings`.

The values you bind to the free variables must be of the type specified by the corresponding entry in the `FreeVariables` map that was specified at query construction. For primitive types, the type of value stored in the `FreeVariableBindings` must be the associated wrapper type. `ObjectStore` does not check that the correct types are bound until it executes the query.

For example, the `INPUT_SALARY` free variable is used in the previous example query. Your application might read in a value from a user in an interactive program or compute the value in some other way. Regardless of how your application computes the value, the free variable is bound to a specific value only when the query is executed. For example:

```
int INPUT_SALARY = {user input or some other computation};
FreeVariableBindings bindings = new FreeVariableBindings();
bindings.put("INPUT_SALARY", new Integer(INPUT_SALARY));
Set result = q.select(employees, bindings);
```

Executing Queries

You can execute a query that

- Specifies predefined variables or free variables
- Returns one element or a set of elements

Obtaining a set

To obtain the set of elements that satisfy a query, call the `com.odi.util.query.Query.select()` method. The two overloadings follow:

```
public Set select(Collection coll)
public Set select(Collection coll,
    FreeVariableBindings freeVariableBindings)
```

The `coll` argument specifies the collection to be queried. If this query has been explicitly optimized with the `Query.optimize()` method, any indexes specified in the optimization must be available on this collection. If this query has not been explicitly optimized, ObjectStore optimizes it for all indexes on the collection being queried. If the query has been explicitly optimized for indexes that are not available on the specified collection, ObjectStore signals `QueryIndexMismatchException`.

The `freeVariableBindings` argument specifies a `FreeVariableBindings` object that defines bindings for each free variable in the query. For each entry, the key is a `String` that identifies the free variable, and the value is the value that should be associated with the free variable during the evaluation of the query. The value must be of the type specified by the corresponding entry in the `FreeVariable` argument passed to the `Query` constructor. For the query to be evaluated, every free variable associated with the query when it was constructed must have a corresponding binding. Also, every free variable binding must correspond to a free variable that was specified when the query was constructed. If the free variable bindings do not match

the free variable definitions specified when the query was constructed, `ObjectStore` signals `QueryException`.

The `select()` method returns a newly allocated transient `Set` that contains the elements that satisfy the query. If `ObjectStore` does not find any matching elements, it returns an empty collection. The returned `Set` is transient.

Obtaining a single element

To obtain one element that satisfies a query, call the `com.odi.util.query.Query.pick()` method. Following are the two overloadings:

```
public Object pick(Collection coll)
public Object pick(Collection coll,
    FreeVariableBindings freeVariableBindings)
```

The `coll` and `freeVariableBindings` arguments are the same as for the `select()` method. The `pick()` methods return the first element found that satisfies the query. If no elements in the collection satisfy the query, `ObjectStore` signals `NoSuchElementException`.

Type of returned element

The `select()` and `pick()` methods never return elements that are not of the class that was specified as the collection element type when the query was constructed.

Null values

Queries ignore null elements but not null fields. The result set of a query never includes null elements. When a query reaches a null element, execution continues to the next element. Suppose you have a query like the following:

```
name != "fred"
```

A query that evaluates this on a collection returns elements with null `name` fields as well as elements with names that are not `"fred"`.

Now suppose you have a query like the following:

```
spouse.name != "fred"
```

On a collection that includes elements that do not have spouses, this query does not return those elements without spouses. It returns only the elements that have spouses with names that are not `"fred"`, plus the elements that have spouses with null `name` fields.

Limitations on Queries

When a query refers to a class or field, the class or field must be public.

When a query refers to a method, the method must return something. In other words, in a query string, you cannot refer to a method that returns void.

Enhancing Query Performance with Indexes

When you want to run a query on a particularly large collection, it is useful to build indexes on the collection to accelerate query processing. An index provides a reverse mapping from a field value, or from the value returned by a method when it is called, to all elements that have the value. A query that refers to an indexed member executes faster because it is not necessary to examine each object in the collection to determine the elements that match the predicate. Also, `ObjectStore` does not need to fetch into memory every element.

This section discusses the following topics:

- How Indexes Work
- Adding Indexes to Collections
- Dropping Indexes from Collections
- Using Multistep Indexes in Queries
- Sample Program That Uses Indexes
- Sample Program That Queries User-Defined Fields
- Modifying Index Values
- Managing Indexes and Index Values
- Optimizing Queries for Indexes
- Manipulating Indexes Outside the Query Facility

How Indexes Work

When you add an index to a collection, `ObjectStore` examines every element of the collection to determine the value of the indexed field or method. After you build the index, you can run queries against the collection without reexamining the elements to determine the values of any indexed members. The query examines the index instead of the collection.

A query can include both indexed fields and methods and nonindexed fields and methods. ObjectStore evaluates the indexed fields and methods first and establishes a preliminary result set. ObjectStore then applies the nonindexed fields and methods to the elements in the preliminary result set.

Adding Indexes to Collections

You can add indexes to any collection that implements the `com.odi.util.IndexedCollection` interface, directly or indirectly. Note that the `IndexedCollection` interface extends the `Collection` interface.

The `IndexedCollection` interface provides methods for adding and removing indexes and updating indexes when the indexed data changes. In this release of ObjectStore, `com.odi.util.OSTreeSet` is the only collection class that already implements `IndexedCollection`. You can, of course, define other `Collection` classes that implement `IndexedCollection`. Call the `com.odi.util.IndexedCollection.addIndex()` method to create an index. Following are the three overloadings:

- `addIndex(Class elementType, String path)`
- `addIndex(Class elementType, String path, boolean ordered, boolean duplicates)`
- `addIndex(Class elementType, String path, boolean ordered, boolean duplicates, Placement placement)`

The `elementType` argument indicates the type to which the index applies. Objects of other types can be in the collection that you index but they are ignored by the index. A query that uses the index does not return such elements.

The `path` argument indicates the member to be indexed. A method member can have no arguments or one constant argument. The path can be either the name of a public field or a call to a public instance method, where the public instance method can be in a superclass. The path can also designate a complex navigation path through multiple public data members, such as `a.b().c.name`. If the syntax is incorrect, ObjectStore signals `IndexException`.

The `ordered` and `duplicates` arguments allow you to specify whether the index is ordered and whether it allows duplicates. If you do not specify the `boolean` arguments, the index is unordered and it allows duplicates.

Finally, the `Placement` parameter indicates the database or segment in which to store the index. If you do not pass a `Placement` argument, `ObjectStore` stores the index in the same database, segment, and cluster as the collection.

Dropping Indexes from Collections

Call the `com.odi.util.IndexedCollection.dropIndex()` method to remove an index from a collection. Following is the method signature:

```
public boolean dropIndex(Class elementType, String path)
```

The `elementType` argument indicates the type to which the index applies.

The `path` argument indicates the member for which the index is being removed.

If the index being dropped is a primary index from an `OSTreeSet`, the method replaces the primary index with a map that uses object hash codes to find the objects in the `OSTreeSet`.

Note Removing an index from a collection destroys the index.

Using Multistep Indexes in Queries

You can create a query that uses a *multistep index*, which is an index on a complex navigational path that accesses multiple public data members. It optimizes queries that use that same path. For example, if you wanted to know all employees whose supervisor has a salary less than 50,000, you could create a multistep index and use it in your query as follows:

```
public class Employee {
    public int salary;
    public Employee supervisor;
    ...
}

// Getting the employees collection
Collection employees = (Collection)db.getRoot("employees");

// Adding the multistep index
employees.addIndex(Employee.class, "supervisor.salary",
    true /*ordered*/, true /*duplicates*/);

...

//Query using the multistep index
Query q=new Query(Employee.class, "supervisor.salary < 50000");
Set result = q.select(employees);
```

Sample Program That Uses Indexes

In the `com/odi/demo/query` directory, the `QueryCustomers` class includes the following example of using an index:

```
IndexedCollection collection = new OSTreeSet(db);
try {
    collection.addIndex(Employee.class, "salary");
} catch (IllegalAccessException e) {
    System.err.println("Couldn't access field: " + e);
    System.exit(1);
}
Set result = q.select(employees);
```

Sample Program That Queries User-Defined Fields

In the `com/odi/demo/props` directory, the generic `PropertiesObject` class allows you to create instances in a database without defining the Java classes for the schema. For more information, see the `README` file.

Modifying Index Values

After you add an index to a collection, `ObjectStore` maintains it automatically as you add or remove elements from the collection. However, it is your responsibility to manage index maintenance when indexed members are modified for instances that are already members of an indexed collection.

For example, suppose you insert `Lee` into your collection of employees. You build an index for this collection on the `phoneExtension` field. A query of `"phoneExtension == 1234"` returns `Lee`. If you remove `Lee` from the collection, `ObjectStore` updates the index so it no longer includes `Lee`. However, if you leave `Lee` in the collection but change `Lee`'s phone extension, you must manually correct the index so that `Lee` refers to the correct phone extension.

Methods

There are three methods that you can use to manually maintain an index:

- `IndexedCollection.removeFromIndex()` removes a value from the index.
- `IndexedCollection.addToIndex()` inserts a value into the index.
- `IndexedCollection.updateIndex()` removes a value from the index and replaces it with a value that you specify.

After an application calls one of these methods, the next time the application uses that index, it uses the updated index. A call to `updateIndex()` does the same thing as a call to `removeIndex()` followed by a call to `addToIndex()`.

The exception is that `removeIndex()` and `addToIndex()` inspect the value to determine the index key. That is, they apply the index's path expression to obtain the key from the value. With `updateIndex()`, you pass in the old key and the new key. `ObjectStore` does not have to inspect the value to determine its key. For this reason, and because there is a single call, using `updateIndex()` is more efficient.

Removing and adding index values

The `removeFromIndex()` method has the following two overloadings:

```
public void removeFromIndex(Object value)
public void removeFromIndex(Class elementType,
    String path, Object value)
```

The `addToIndex()` method has the following two parallel overloadings:

```
public void addToIndex(Object value)
public void addToIndex(Class elementType,
    String path, Object value)
```

Usually, after you remove a value from an index, you should add a value to replace it.

If you know exactly the value that you need to add or remove, you can use the form that specifies `elementType`, `path`, and `value`. If you do not know the indexes that exist, or if you modified a lot of different fields and want to update all indexes, use the short form. In this case, `ObjectStore` iterates over all indexes and updates all of them.

Following is an example of removing and adding values to an index:

```
Employee lee = new Employee("Lee", 1234);
collection.insert(lee);
try {
    collection.removeFromIndex(lee);
    lee.setExtension(5678);
    collection.addToIndex(lee);
} catch (IllegalAccessException e) {
    System.err.println("Could not access field: " + e);
    System.exit(1);
}
```

Updating indexes

The `updateIndex()` method has the following signature:

```
public void updateIndex(Class elementType,
    String path, Object oldKey, Object newKey, Object value)
```

Following is an example of updating an index:

```
Employee lee = new Employee("Lee", 1234);
collection.insert(lee);
```

```
lee.setExtension(5678);
collection.updateIndex(
    Employee.class, "extension",
    new Integer(1234), new Integer(5678), lee);
```

Managing Indexes and Index Values

When you add or drop an index, you do it at the class level. That is, you specify the class and member that the index is on. For example, you might add an index on the `name` field of the `Employee` class, as follows:

```
employeeCollection.addIndex(Employee, "name")
```

However, when you perform maintenance on an index, that is, when you call `removeFromIndex()`, `addToIndex()`, or `updateIndex()`, you do it at the instance level. For example, suppose you have an employee named Jones with an employee ID number of 1234. The employee's name changes to Smith. You must update this index entry at the instance level by

- 1 Removing the object from the index while the object still has its old key value
- 2 Adding the object back into the index with the new key value

The following example shows how to update the index when an employee's name changes from Jones to Smith:

```
employeeCollection.removeFromIndex(employee1234);
employee1234.setName("Smith");
employeeCollection.addToIndex(employee1234);
```

For each index on the `Employee` class, these methods update the index's value for `employee1234`. If there are multiple indexes on `Employee`, the one-argument overloading of `removeFromIndex()` and `addToIndex()` updates all of them. You need not specify that you want to update the index on the `name` field. For example, there might be indexes on the `Employee.salary` and `Employee.location` fields, as well as the `Employee.name` field. The previous code fragment would update the indexes on `salary` and `location`, as well as the index on `name`, even though only the index on `name` needs to be updated. This technique is useful when you make a lot of changes to different fields.

If you use the three-argument overloading of `removeFromIndex()` or `addToIndex()`, you can update only the index that needs to be updated. You must know the type of the indexed element, the name of the indexed member, and the value to be removed or added. For example:

```
employeeCollection.removeFromIndex(
```

```

    Employee.class, "name", employee1234);
employee1234.setName("Smith");
employeeCollection.addToIndex(
    Employee.class, "name", employee1234);

```

You can also update an index by using the `updateIndex` method. Instead of removing the object with its old key value from the index and then adding the object with its new key value back into the index, you can just call the `updateIndex` method.

If you call the `updateIndex` method by using the previous example, and then your code would look like the following:

```

employeeCollection.updateIndex(
    Employee.class, "name", employee1234.name, "Smith",
    employee1234)
employee1234.setName("Smith");

```

Optimizing Queries for Indexes

If you do not explicitly optimize a query for a particular set of indexes, `ObjectStore` optimizes the query automatically when it applies the query to a collection. This means that `ObjectStore` optimizes the query to use exactly those indexes that are available on the collection being queried.

Preparation Before you optimize a query, you must obtain an instance of `IndexDescriptorSet`. An `IndexDescriptorSet` implements a set of `IndexDescriptor` objects. An `IndexDescriptor` is an object that describes an `IndexMap` on an instance of `IndexedCollection`. Typically, you can obtain an `IndexDescriptorSet` with a call to `IndexedCollection.getIndexes()` on any collection that has exactly the indexes for which you want to optimize your query.

Explicit optimization To explicitly optimize a query, call the `Query.optimize()` method. The method signature is

```
public synchronized void optimize(IndexDescriptorSet indexes)
```

The `indexes` argument is an instance of `IndexDescriptorSet` that contains `IndexDescriptor` objects that describe the indexes against which to optimize.

Reoptimizing If you apply an optimized query to the same collection again or to another collection with the same indexes, `ObjectStore` uses the same optimization. Reoptimization is not required. However, suppose you apply an optimized query to a collection that does not have all the indexes that were present when the query was first run. In this situation, `ObjectStore` must reoptimize

the query. `ObjectStore` does this automatically; your intervention is not required.

Manual optimization

Automatic index optimization is convenient and effective. However, suppose a query is to be run multiple times against more than one collection, potentially with different indexes available. In this situation, it might be best to manually control the query optimization strategy.

For example, consider that the same query is to be run repeatedly against two different collections, in which the collections have different indexes. One alternative is to create two separate query objects, one for each collection. This avoids the overhead of recomputing the indexing optimization strategy each time you apply the query to a different collection. A second alternative is to explicitly optimize a query to use only the intersection of the indexes that are available on both collections. You can do this with a call to `Query.optimize()`. Pass in an `IndexDescriptorSet` object that contains descriptions of only the common indexes.

Restriction

If you explicitly optimize a query with the `Query.optimize()` method, it cannot run against a collection that does not have the specified indexes. If you try to do this, `ObjectStore` signals `QueryIndexMismatchException`. In this way, an explicitly optimized query differs from an automatically optimized query. An automatically optimized query reoptimizes itself as needed when you run it against a collection with different indexes.

This might be useful when it would be undesirable to run a particular query on a collection that does not have the required indexes. For example, this is useful when the collection is very large and the overhead of examining every element of the collection is prohibitive.

To evaluate query expressions efficiently, `ObjectStore` compiles query expressions into classes and methods that are loaded when the query is evaluated. Each new query potentially can result in the creation of a new class with a new internal name to represent the compiled state of the query. When the query is no longer referenced, this class is normally garbage collected by the Java VM GC and its storage reclaimed.

-Xnoclassgc option to Java VM

If you are using JDK 1.2 with the `-Xnoclassgc` option to the Java VM, you risk running out of heap storage as the query expression classes that are generated by `ObjectStore` accumulate over time. The `-Xnoclassgc` option prevents the query expression classes from being garbage collected.

Manipulating Indexes Outside the Query Facility

You can use the `IndexMap` interface to directly access and manipulate indexes outside the query facility. This interface is useful when you want a sorted result set and you can represent the query as a single range expression on an indexed member. Instead of running a query, you can iterate over the index directly. See the on-line *Java API Reference* for more information on `com.odi.util.IndexMap`.

Storing Objects as Keys in Persistent Hash Tables

The `com.odi.util.OMHashtable` class introduces a new requirement for classes of objects that will be stored as keys in persistent collections: these classes must provide a suitable `hashCode()` method. `ObjectStore` and the class file postprocessor provide facilities for doing this conveniently.

This section discusses the following topics:

- Requirements for Hash Code Methods
- Providing an Appropriate Persistent Hash Code Method
- Storing Built-In Types as Keys in Persistent Hash Tables

Requirements for Hash Code Methods

Objects that are stored as keys in persistent hash tables must provide hash codes that remain the same across transactions. `ObjectStore` can create a new transient Java object in each transaction to represent a particular persistent object, so it is important that the `hashCode()` method used for persistent objects return the same hash code for these different transient objects.

The default `Object.hashCode()` method supplies an identity-based hash code. This identity hash code might depend on the virtual memory address or some internal implementation-level metadata associated with the object. Such a hash code is unsuitable for use in a persistent identity-based hash table because it would effectively be different each time an object was fetched in a transaction.

Providing an Appropriate Persistent Hash Code Method

In cases in which a persistence-capable class does not override the `hashCode()` method it inherits from `Object`, the class file postprocessor arranges for the class to implement a `hashCode()` method suitable for storing instances in persistent hash tables. It does this by adding an `int` field to the class. This field is initialized to an appropriate hash code when an instance is created and returns the value stored in the field from its `hashCode()` method. This hash code value is guaranteed to remain unchanged for the lifetime of the object.

Applications need to provide their own `hashCode()` methods for classes that define `equals()` methods that depend on the contents of instances rather than on object identity. If the `equals()` method only uses the `==` operator to compare the argument with `this` (or inherits `Object.equals()`), it is identity based and the `hashCode()` method provided by the class file postprocessor is appropriate. If the `equals()` method compares the contents of the objects, it is contents based and your application must supply a `hashCode()` method that returns the same hash code value for all objects whose contents make them return `true` when compared to the `equals()` method.

If an application does not need to store instances of a particular persistence-capable class as keys in a persistent hash table, there is no special requirement for that class's `hashCode()` method. In this case, to avoid making all your instances one word larger, have the class define or inherit a `hashCode()` method that calls the superclass's `hashCode()` method:

```
public int hashCode() {return super.hashCode();}
```

Doing this ensures that the `hashCode()` method inherited from `Object` is used, which returns a hash code that can be used only in a nonpersistent context.

Storing Built-In Types as Keys in Persistent Hash Tables

You can use the following built-in Java types as `OSHashtable` keys without overriding the `hashCode()` method:

- `java.lang.String`
- Wrapper classes, for example, `Character`, `Integer`, `Long`, `Float`

There is no way to override the `hashCode()` method for arrays. Therefore, do not use Java arrays as keys in persistent hash tables. You can, however,

define a class that stores the array as a field and provides an appropriate `hashCode()` method.

Java wrapper classes work well as keys because their `hashCode()` methods are based on the value of the object rather than on its address.

Using Third-Party Collections Libraries

You can use a third-party Java collections library with `ObjectStore`. The advantages of doing so are that it might have features that you need or that you might be familiar with its use. The disadvantage is that it might not scale to the degree that you need.

One third-party library you can use is Doug Lea's collections library. An example of using this is in the `collection` subdirectory of the `ObjectStore` `demo` directory.

Chapter 8

Generating Persistence-Capable Classes Automatically

This chapter provides information and instructions for using the class file postprocessor to make classes persistence capable. Reference information for all postprocessor options is in Chapter 14, Tools Reference, on page 323. For information on what Java-supplied classes are persistence capable, see Java-Supplied Persistence-Capable Classes on page 311.

Caution

For simple applications, it is best to postprocess all classes together. For more complex applications, you can postprocess your classes in correctly grouped batches. See *Postprocessing a Batch of Files Is Important* on page 198.

Failure to postprocess the correct classes together can result in problem situations that appear when you try to run the application and that are difficult to diagnose. There are postprocessor options that allow you to determine those classes that are made persistence capable.

Contents

This chapter discusses the following topics:

Overview of the Class File Postprocessor	196
Running the Postprocessor	200
Managing Annotated Class Files	208
Creating Persistence-Aware Classes	212
How the Postprocessor Works	213
Including Transient and Already Annotated Classes	219
Putting Processed Classes in a New Package	221

Creating Persistence-Capable Classes with Transient Fields	224
Customizing Updated Classes	226
Optimizing Operations That Retrieve Persistent Objects	230
Performing a Test Run of the Postprocessor	232
Using an Input File	233
Annotations You Must Add	234
Class File Postprocessor Limitations	236

Overview of the Class File Postprocessor

To store an object in a database, the object must be persistence capable. For an object to be persistence capable, it must include code that allows persistence. ObjectStore includes the class file postprocessor utility to insert the required code, referred to as *annotations*, into your class files automatically.

Annotating classes automatically

The command you use to run the class file postprocessor command-line utility is `osjcfp`. The postprocessor provides a number of command options that allow you to tailor the results to your needs.

You can run the postprocessor (or its companion API) on classes or class libraries that you create or that you purchase from a vendor. See com/odi/demo/collections/README.htm for an example of making a third-party library persistence capable.

Note

You must explicitly postprocess each class that you want to be persistence capable by using the `osjcfp` utility. When you extend a persistence-capable class, objects do not inherit persistence, which is the ability to be stored in a database.

When you postprocess or manually annotate a class, this registers the class with ObjectStore. If a class is not postprocessed or manually annotated, ObjectStore signals `ClassNotRegisteredException`.

This overview provides the following information:

- Description of the Annotations
- Description of the Process
- Postprocessing a Batch of Files Is Important
- Manual Annotation

Description of the Annotations

The class file postprocessor annotates classes you define so that they are persistence capable. This means that the postprocessor makes a copy of your class files, places them in a directory you specify (either the source directory or another directory), and adds byte-code instructions (annotations) that are required for persistence.

These annotations are

- Modifying the class to implement the `com.odi.IPersistent` interface.
- Defining methods to initialize instance fields with data from the database, writing modified fields to the database, and resetting instance fields to default values.
- Modifying methods to fetch the contents of persistent instances from the database as needed and to mark modified instances so their changes can be written to the database at transaction commit.

Before an application can access the contents of a persistent object, it must call the `ObjectStore.fetch()` method to read the object or the `ObjectStore.dirty()` method to modify the object. These calls make the contents of the object available to your application. The postprocessor inserts these calls in methods of classes that it makes persistence capable or persistence aware.

- Defining an additional class that provides schema information about the persistence-capable class, if specified or required. This new class is a subclass of the `com.odi.ClassInfo` class.

Description of the Process

Before you run the postprocessor, you must compile your source files. The set of files you run the postprocessor on can contain a combination of class files, `.zip` files, and `.jar` files. The postprocessor generates annotated class files and places them in a directory that you specify.

This destination directory is never the original directory unless you specify the `-inplace` option (see page 326). When you are in a development cycle, it is best to specify a directory other than the original directory. Doing so avoids errors and provides both a persistence-capable and a transient version of the same class.

It is not necessary to recompile all classes before iteratively running the postprocessor. The requirement is that the compiled classes be consistent.

The postprocessor tries to minimize the amount of work it does. It checks file modification times and reprocesses only those files that have changed.

Postprocessing a Batch of Files Is Important

In one execution of the postprocessor, the postprocessor must operate on a correctly grouped set of files. For example, an application might use a file, perhaps a library, that is already annotated. You must not specify the annotated files when you run the postprocessor on the rest of the files in your application. Hence, the term *batch* means all files that the postprocessor must annotate in one execution of the `osjcfp` command. Each batch must have its own postprocessor destination directory for this to work correctly.

You can use the postprocessor `-inplace` option to create multiple batches. When you do, there is no requirement for the separate batches to be stored in different directories.

Example of one batch

When you write a program that uses persistence, the program usually consists of a batch (a set) of classes, for example, classes `A`, `B`, and `C`. They typically are defined in files called `A.java`, `B.java`, and `C.java`. It is possible for each class to reference the other classes. For example, `B` might refer to `C`, and `C` might refer to `B`. There is no ordering or layering; there are no rules for references among the classes.

When this is the scenario, you must run the postprocessor on all of these classes at the same time. You cannot run the postprocessor on each file individually. This is because when the postprocessor operates on `A`, it might refer to `B` and `C`. The postprocessor must have information about `B` and `C` to correctly annotate `A`.

Example of two batches

In relatively simple programs, there is only one batch involved. However, sometimes there might be more than one batch in an application. Suppose, for example, that you want to write a persistent program that uses an existing library. An example of this is `djg1`, which is the ObjectStore persistence-capable version of ObjectSpace's JGL library. Your program consists of `A`, `B`, and `C` plus the JGL library.

In a simple (one-batch) program, when you run the postprocessor, you always specify all files in your application. In this case, you do not want the postprocessor to operate on JGL because it has already been postprocessed. In fact, you probably do not have the class files that have not been postprocessed.

It is correct to run the postprocessor on only `A`, `B`, and `C`. This is because there is a rule: JGL classes never know about `A`, `B`, and `C`. After all, JGL was written, finished, and put on the shelf before `A`, `B`, and `C` were created.

There are two batches here:

- The first batch contains the persistence-capable JGL library. `ObjectStore` runs the postprocessor on this batch.
- The second batch contains your own classes, `A`, `B`, and `C`. You run the postprocessor on this batch.

Whenever you run the postprocessor, you must run it on a whole batch. Each batch must have its own postprocessor destination directory.

Checking for
correct
batches

To determine whether you have correctly grouped your files in batches, you can apply this rule: Class `A` and class `B` must be in the same batch if either of the following is true:

- Class `B` inherits from class `A` and either class is persistence capable.
- Class `A` is persistence capable or persistence aware and it refers directly to the fields of class `B`, which is persistence capable.

Postprocessor API

`ObjectStore` also includes a companion API for its class file postprocessor utility for those times when it is useful to call the postprocessor from your Java application. The method that you call is

```
com.odi.filter.OSCFP.filter(String[] argv)
```

For information on how to use this method, see the Postprocessor API on page 332.

Manual Annotation

In exceptional situations, you might want to insert all required annotations needed for persistence and not use the postprocessor at all. For more information, see Chapter 9, *Generating Persistence-Capable Classes Manually*, on page 237. You can also manually annotate your code to meet some persistence requirements and then run the postprocessor to insert the other annotations.

Running the Postprocessor

To make classes persistence capable, do the following:

- 1 Compile the source files.
- 2 Run the postprocessor on the resulting class files.

You must run the postprocessor on all class files in a batch at the same time.

Some Java-supplied classes are persistence capable. Others are not persistence capable and cannot be made persistence capable. A third category of classes can be made persistence capable but there are important issues to consider when you do so. Be sure to read *Java-Supplied Persistence-Capable Classes* on page 311.

The topics discussed in this section are

- Preparing to Run the Postprocessor
- Requirements for Running the Postprocessor
- Example of Running the Postprocessor
- About the Postprocessor Destination Directory
- How the Postprocessor Interprets File Names
- Order of Processing
- How the Postprocessor Handles Duplicate File Specifications
- How the Postprocessor Handles Files Not Found
- .Zip and .Jar Files as Input to the Postprocessor
- How the Postprocessor Handles Previously Annotated Classes
- Troubleshooting OutOfMemory Error
- How the Postprocessor Handles Inner Classes
- When ClassInfo.java Files Are Generated

Preparing to Run the Postprocessor

Before you run the postprocessor, ensure that the following `.jar` files are explicitly specified in your `CLASSPATH`. An entry for the directory containing them is not sufficient.

- A `tools.jar` entry must be in your `CLASSPATH` environment variable.
- The `osji.jar` file must be in your `CLASSPATH` environment variable.

- Also, you must update your `PATH` environment variable to contain the `bin` directory from the ObjectStore for Java distribution.

On Windows, you might set `PATH` to be something such as the following:

```
PATH=c:\odi\ostore\bin;c:\odi\osji\bin;  
c:\winnt\system32;c:\winnt
```

On UNIX, it would be something such as the following:

```
PATH=/usr/ucb:/usr/bin:/opt/SUNWspro/bin:  
/opt/ODI/ostore/bin:/opt/ODI/osji/bin
```

Requirements for Running the Postprocessor

The postprocessor requires specification of

- The `-dest` option with a destination directory for the annotated class files. This can be an absolute or a relative path name. This directory must already exist when you specify it on the postprocessor command line. The postprocessor does not create it.

You can also specify the `-inplace` option to instruct the postprocessor to overwrite your original class files. If you do, the `-dest` option is still required.

- A batch of files. A batch includes all files in your application except already annotated files that your application refers to. You can specify
 - One or more `.class` files
 - One or more `.zip` files
 - One or more `.jar` files
 - One or more class names (you must specify the package names)
 - Any combination of the previous items

Insert a space between specifications and be sure to specify the required destination parameter. When you run the postprocessor, each batch must have its own destination directory. For example:

```
osjcfp -dest osjcfpout com.odi.demo.threads.Institution Banking.jar  
Account.class
```

You can specify additional options, which are described in `osjcfp: Running the Postprocessor` on page 323.

Example of Running the Postprocessor

To make the `Person` class persistence capable, enter a command such as the following:

```
osjcfp -dest ..\osjcfpout Person.class
```

This command assumes that the `Person.class` file is in the current directory and the `osjcfpout` directory is a sibling to the current directory. When the postprocessor successfully generates an annotated version of the `Person.class`, the file contains the information `ObjectStore` needs to store instances of `Person` persistently.

The postprocessor places the annotated `Person.class` file in a package-relative subdirectory of the `osjcfpout` directory. For example, suppose the `Person` class package name is `com.odi.demo.people`. Further suppose that the `osjcfpout` directory is in the `\users\kim` directory. The postprocessor writes the annotated class file to a file whose name is made up of the destination directory, the class package, the class name, and the `.class` extension:

```
\users\kim\osjcfpout\comcom\odi\demo\people\Person.class
```

Note that both of the following commands have the same results, as specified previously:

```
osjcfp -dest ..\osjcfpout Person.class
osjcfp -dest \users\kim\osjcfpout com.odi.demo.people.Person
```

About the Postprocessor Destination Directory

The postprocessor never overwrites the class files specified on the postprocessor command line unless you specify the `-inplace` option when you run the postprocessor. If you do specify the `-inplace` option, the postprocessor does overwrite the original class files with the annotated class files.

If you specify a destination directory in such a way that it would store the annotated class file in the same location as the unannotated class file and you do not specify the `-inplace` option, the postprocessor displays an error message and terminates. It does not produce any class file output.

The postprocessor ignores classes that are rooted in the destination directory. If you try to postprocess a class that exists only in the destination directory and you do not specify `-inplace`, the postprocessor reports that it cannot find the file. For example, if you specify the following command when you run `osjcfp`, you receive an error as shown:

```
setenv CLASSPATH
/usr/devo/java/test:/opt/ODI/pse.jarosji.jar:/opt/ODI/tools.jar
cd /usr/devo/java/test
javac com/users/jobs/teacher.java
osjcfp -d . com.users.jobs.teacher
```

Error: Class com/users/jobs/teacher could not be found.

Because the postprocessor ignores the destination directory in the `CLASSPATH` when it looks up classes, it is unable to locate the specified class. Consequently, the destination directory you specify cannot be the root directory for any of the classes you want to postprocess or any classes referenced by classes you want to postprocess.

Typically, after you run the postprocessor, you have a transient version of a class (your original file) and a persistence-capable version of the class (in the destination directory).

If there are no errors, the postprocessor places a version of all files specified on the command line in the destination directory. The postprocessor annotates those files that require annotations and does not modify those files that do not.

How the Postprocessor Interprets File Names

If a name you specify ends with `.class`, `.zip`, or `.jar`, the postprocessor assumes that it is an explicit file name for a class file, `.zip` file, or `.jar` file, respectively.

If a name you specify does not end with `.class`, `.zip`, or `.jar`, the postprocessor assumes that it is a class name delimited with periods, for example, `a.b.C`. The postprocessor uses the `CLASSPATH` environment variable or the `-classpath` specification on the postprocessor command line to locate the `.class` file, which can be in a `.zip` file or `.jar` file. (The use of the `-classpath` option does not affect the class path used for the execution of the postprocessor.)

Following is an example of adding the `-classpath` option. It assumes that you are using ObjectStore on UNIX with the JDK installed in `/usr/local/JDK-1.2`.

```
osjcfp -dest osjcfpout -classpath \
/usr/osji:/usr/local/JDK-1.2/java/lib/classes.jar:\
/usr/osji/lib/osji.jar com.odi.demo.threads.Institution \
Banking.jar Account.class
```

CLASSPATH and -classpath The postprocessor uses the class path you specify in the command line to locate the specified files. This is in place of the `CLASSPATH` environment variable. At run time, Java implementations append the location of the system classes to the end of the `CLASSPATH` environment variable. You must do this manually if you specify the `-classpath` option. This is shown in the previous examples as `classes.jar`.

Order of Processing

The postprocessor processes the class files in the order in which they appear on the command line and according to the persistence mode that is in effect when the postprocessor reaches the file name. The persistence mode indicates whether the postprocessor is

- Annotating the class to be persistence capable
- Annotating the class to be persistence aware
- Copying the class to the destination directory without annotating it

Persistence mode options The default persistence mode is that the postprocessor generates persistence-capable classes. Following are the options you can specify to determine the persistence mode:

<i>Persistence Mode Option</i>	<i>Description</i>
<code>-pc</code> <code>-persistcapable</code>	Classes specified after this option are made persistence capable. The postprocessor annotates these classes to include all annotations required by <code>ObjectStore</code> for an object to be persistent. This is the default. If you do not specify a persistence mode option in the postprocessor command line, the postprocessor makes all specified classes and superclasses of those classes persistence capable.

<i>Persistence Mode Option</i>	<i>Description</i>
<code>-pa</code> <code>-persistaware</code>	Classes specified after this option are made persistence aware. The postprocessor annotates these classes so that they can operate on persistent objects, but instances of these classes cannot themselves be stored persistently.
<code>-cc</code> <code>-copyclass</code>	Classes specified after this option are not annotated. Specify this option for classes that should not be annotated either because they are nonpersistent or are already annotated. The postprocessor copies these classes to the destination directory along with the annotated classes.

If you specify a `.class` file or class name, the postprocessor processes it according to the persistence mode that is in effect when the postprocessor reaches the file name. If you specify a `.zip` file or `.jar` file, the postprocessor processes all class files in the `.zip` file or `.jar` file according to the persistence mode that is in effect when the postprocessor reaches the name of the `.zip` file or `.jar` file in the command line. For example:

```
osjcfp -dest osjcfpout -persistaware Tent.class Family.class \
-persistcapable Campers.jar Site.class -copyclass Weather.class
```

Example

After you run the postprocessor with the previous command,

- The `Tent` and `Family` classes are persistence aware.
- The `Site` class, its superclass if it has one other than `java.lang.Object`, all classes in the `Campers.jar` file, and any of their superclasses (other than `java.lang.Object`) are persistence capable.
- The `Weather` class is not annotated and is copied as it is to the destination directory.

How the Postprocessor Handles Duplicate File Specifications

It is permissible for a class to be specified more than once in a command line. For example, a file can be in a `.zip` file and you can also explicitly specify it. On UNIX and on Windows NT using a Sun JDK, a file can be included in a wildcard specification and you can also explicitly specify it. In the previous

example, the `Family` class could be in the `Campers.jar` file. If it were, the postprocessor would annotate the `Family` class to be persistence capable. This is because making a class persistence capable supersedes making it persistence aware. Likewise, making a class persistence aware supersedes copying it as is to the destination directory.

If you specify the same class more than once on a command line, both specifications must resolve to the same disk location. For example, suppose you specify both `Person.class` and `com.odi.demo.people.Person`. This is allowed only if the class path causes `com.odi.demo.people.Person` to resolve to the same `Person.class` that is explicitly specified.

How the Postprocessor Handles Files Not Found

The postprocessor must be able to find every file that you specify on the command line. If it cannot find one or more files, it displays an error message and stops processing. It does not produce any annotated class files.

.Zip and .Jar Files as Input to the Postprocessor

If a class originates in a `.zip` file or `.jar` file, either because you specify a `.zip` file or `.jar` file when you run the postprocessor or because the class path search locates the class in a `.zip` file or `.jar` file, the postprocessor writes the annotated class to the package-appropriate subdirectory of the destination directory.

How the Postprocessor Handles Previously Annotated Classes

If the postprocessor previously annotated a `.class` file, you can specify only that `.class` file to be copied. You cannot specify it to be annotated. If you do, the postprocessor displays a message that states the specified class that was already annotated and terminates without producing any annotated files.

Troubleshooting OutOfMemory Error

The JDK 1.2 imposes a memory limitation of 16 MB unless you override it. If you receive a `java.lang.OutOfMemory` error during postprocessing, you must increase the run-time memory pool. Do one of the following:

- Set the `OSJCFPJAVA` environment variable to include the `-Xmx` option. For example, Solaris `csh` users can enter

```
setenv OSJCFPJAVA "java -Xmx32m"
```

Windows users can enter

```
set OSJCFPJAVA=java -Xmx32m
```

- Edit the `osjcfp` script (Solaris) or `osjcfp.bat` script (Windows) to incorporate the `-Xmx` option in the invocation of Java near the end of the script. On Solaris, the line to change is

```
$OSJCFPJAVA $javaargs com.odi.filter.OSCFP $args
```

On Windows, the line to change is

```
%osjcfpjava% com.odi.filter.OSCFP %1 %2 %3 %4 %5 %6 %7 %8
```

Add `-Xmx32m` before the `com.odi.filter.OSCFP` entry. This allows the Java virtual machine to increase the heap to 32 MB. You can increase this value further if you need to.

How the Postprocessor Handles Inner Classes

When you define a class inside another class, you must explicitly make both the outer class and the inner class persistence-capable. For example, suppose you define the following class:

```
Class Foo {
    int a;
    public class Bar {}
}
```

You must specify both the `Foo` class and the `Bar` class when you run the postprocessor:

```
osjcfp -dest ../osjcfpout -pc Foo.class -pc Foo$Bar.class
```

When ClassInfo.java Files Are Generated

The postprocessor always generates `xxxClassInfo` classes for persistence-capable classes that are private. However, by default, the postprocessor does not generate `xxxClassInfo` classes for persistence-capable classes that are public. Instead, the postprocessor relies on the ability of ObjectStore to build an `xxxClassInfo` class dynamically when needed, using the reflection API. This optimization reduces the disk footprint and application start-up times because there are fewer classes to load when the application starts.

The reflection API is subject to security and access constraints that are enforced to varying degrees at run time, depending on the version of your Java Virtual Machine and your platform. If your application encounters run-time security errors while attempting to generate `xxxClassInfo` classes dynamically, specify the `-nooptimizeclassinfo` option when you run the

postprocessor. When you specify `-nooptimizeclassinfo`, the postprocessor generates the `xxxClassInfo` classes; therefore, the reflection API is not used.

Managing Annotated Class Files

After you run the postprocessor, there are two versions of your class files:

- The unannotated class files in the original directory
- The annotated class files in the destination directory

It is important to keep these versions separate because

- When you compile source code, you must ensure that any class files the compiler reads in are unannotated class files. The compiler must find unannotated class files before it finds annotated class files with the same names.

Note: While the above is true for simple applications, it might not be true for more complex applications. See *Using the Right Class Files in Complex Applications* on page 210.

- When you run your application, you must ensure that `ObjectStore` finds the annotated class files before it finds the unannotated class files with the same names.

There are several ways to accomplish this. Technical Support recommends that you

- Specify the `-classpath` argument to the compiler so it finds the unannotated class files first
- Modify your `CLASSPATH` environment variable so Java can find the annotated class files first when it runs your application

To help you manage annotated class files, this section discusses

- Ensuring That the Compiler Finds Unannotated Class Files
- Ensuring That `ObjectStore` Finds Annotated Class Files
- Using the Right Class Files in Complex Applications
- Alternatives for Finding the Right Files
- How the Postprocessor Determines Whether to Generate an Annotated Class File

Ensuring That the Compiler Finds Unannotated Class Files

The compiler can locate class files in the following two ways:

- Through the `CLASSPATH` environment variable
- Through the `-classpath` argument to the compiler

`CLASSPATH` is convenient when you compile, but when you try to run your application, `ObjectStore` finds the unannotated files before it finds the annotated files. The `-classpath` option is more cumbersome to use because it means that the path to Java system classes must be listed explicitly in the argument, but it is safe. It ensures that the compiler does not operate on annotated class files.

Example 1

For example, suppose that `ObjectStore` is installed in `c:\osji` and you are building an application in `c:\app`. Your destination directory for annotated class files is `c:\app\osjcfpout`. Your `CLASSPATH` variable might look like the following:

```
CLASSPATH=c:\osji\osji.jar;c:\app\osjcfpout;c:\app
```

When you run the compiler, specify the `-classpath` option with the following path. This removes the destination directory from the class look-up path and adds the Java classes to the path.

```
javac -classpath c:\app;c:\osji\osji.jar;c:\jdk12\lib\classes.jar App.java
```

Example 2

Following is an example of why it is important for the compiler to operate on unannotated class files. Suppose you have two classes named `X` and `Y` in the same postprocessor batch. Neither of these classes is explicitly declared to implement `com.odi.IPersistent`. Now suppose you add the following two methods to class `Y`:

```
void foo(com.odi.IPersistent p) {}
void bar() { foo(new X()); } // Trying to pass an X instance to
    // a function that is expecting com.odi.IPersistent
```

If you recompile only `Y.java` and the compiler finds the annotated classes, examination of the annotated class file allows the compiler to determine that `X` implements `IPersistent`, which allows `Y.bar()` to compile. If you then recompile both `X` and `Y`, the compiler recognizes that `X` is not declared to implement `com.odi.IPersistent` and refuses to compile class `Y`, even though it compiled successfully earlier.

Ensuring That ObjectStore Finds Annotated Class Files

When you run your application, ObjectStore must find the annotated class files before it finds the unannotated class files. The recommended way to ensure this is to define a `CLASSPATH` environment variable that has the postprocessor destination directory before the source file directory.

Example

Consider the following example:

- ObjectStore is installed in `c:\odi\osji`.
- You are building an application in `c:\app`.
- You create a directory named `c:\app\osjcfpout` to hold the annotated class files.

In this scenario, use the following `CLASSPATH`:

```
c:\app\osjcfpout;c:\app;c:\odi\osji\osji.jar
```

After you modify your `CLASSPATH` environment variable, you can run the postprocessor with no special action. The postprocessor excludes the destination directory from the class path when it does class-path-based searches.

Using the Right Class Files in Complex Applications

There are situations in which you want the compiler to read in annotated class files. In these cases, the referenced classes are similar to an independent library on which you are building your application. The referenced classes form a batch, which is a group of class files that must be postprocessed together. The other files in your application form a second batch.

Independent library

For example, suppose this second batch is named `x`. Specify the `-classpath` option so that it points to the

- Unannotated class files for any classes in `x`
- Annotated class files for any classes that are in other batches and are referenced by classes in `x`

This is the most common multiple-batch scenario. Your application is in one batch and the other batches are existing reusable libraries. Each batch has its own postprocessor destination directory.

Classes referenced by other classes

Now suppose that you are not using an existing library. Your application itself contains a group of referenced classes (first batch), then another group of classes (second batch) that reference the first batch. The following instructions show how to build your application in stages:

- 1 Compile the source files and postprocess the class files in the first batch. (This is the batch of files that are referenced by other classes in the application.)
- 2 Compile the source files in the second batch. You might not want to compile all files in this batch at the same time. Specify the `-classpath` option to point to the annotated class files in the first batch and any unannotated class files in the second batch.
- 3 Run the postprocessor on the second batch. Specify a destination directory that is different from the destination directory that was specified when the postprocessor operated on the first batch. You can package the result of postprocessing the second batch in a `.zip` file or `.jar` file.

Alternatives for Finding the Right Files

In some circumstances, updating your `CLASSPATH` environment variable might be cumbersome or might not work well with your development environment. (This is true for the Symantec Cafe product.) In these cases, you can copy annotated files back to the building directory. However, if you do this, you must remove the annotated files before you recompile. This ensures that subsequent compilation and postprocessing operate on unannotated class files.

Two other alternatives are to

- Delete the contents of the destination directory before you recompile
- Specify the `-classpath` option when you run the postprocessor, just as you did for compilation

How the Postprocessor Determines Whether to Generate an Annotated Class File

When you run the postprocessor, it checks whether any annotated file it is going to create already exists. If an annotated file does not already exist, the postprocessor generates it. If an annotated file does exist, the postprocessor compares the date on the compiled input file with the date on the annotated output file. If the input file date is after the output file date, the postprocessor generates a new output file. If the input file date is before the output file date, the postprocessor does not generate a new file. It assumes that the annotated file that already exists is still valid.

This works well when you run the postprocessor repeatedly with the same command line. However, when you change input parameters to the

postprocessor, it is a good idea to remove the previously annotated class files from the destination directory. The reason for this is that a comparison of dates might not cause a new annotated file to be generated when the specification of a new input parameter requires a new annotated file to be generated.

To force the postprocessor to overwrite existing annotated files, specify the `-f` or `-force` option when you run the postprocessor.

Creating Persistence-Aware Classes

If you know that a class will never need to be stored persistently, you can run the postprocessor to make the class persistence aware. A persistence-aware class can operate on persistent objects but cannot be persistent itself. For an example of how you might use persistence-aware classes, see

com/odi/demo/pport/README.htm.

Persistence-aware annotations require less space than persistence-capable annotations require. The postprocessor adds calls to `ObjectStore.fetch()` and `ObjectStore.dirty()` only where they are needed to operate on persistent objects. When the postprocessor makes a class persistence aware, it does not annotate that class's superclass. You need only make a class persistence aware, instead of copying it as is, if

- The class accesses fields of a persistence-capable class instead of using methods to access the fields.
- The class accesses elements of persistent arrays.

You must make a class persistence aware (or persistence capable) when it includes methods that obtain arrays from persistent objects.

Specifying the Postprocessor Command Line

To create a persistence-aware class, specify the `-pa` or `-persistaware` option followed by the names of the classes that you want to be persistence aware. For example:

```
osjcfp -dest osjcfpout -persistaware Compute.class
```

The preceding command line annotates `Compute.class` so that it has calls to the `fetch()` and `dirty()` methods.

No Changes to Superclasses

Another reason to make a class persistence aware is that doing so does not require changing its superclasses. This is important for classes such as `java.lang.Thread`, whose superclass should not be modified. `java.lang.Thread` is inherently transient, so it makes no sense for it to become persistent because it is not useful when you take it out of the database. Typically, Java system classes are restricted from annotations by the postprocessor.

How the Postprocessor Works

This section describes postprocessor behavior relative to various components in your application. It is important to be familiar with the information here so that the postprocessor produces the results you expect. The topics covered in this section are

- Ensuring Consistent Class Files
- Modifications to Superclasses
- Effects on Inheritance
- Location of Annotated Class Files
- Postprocessor Errors and Warnings
- Handling of final Fields
- Handling of Static Fields
- Which Java Executable to Use
- Line-Number and Local-Variable Information
- Using a Debugger
- Handling of `finalize()` Methods
- Description of Postprocessor Optimizations

Ensuring Consistent Class Files

When you run the postprocessor on more than one class file at a time, all specified classes must be consistent. To ensure class consistency, compile all classes together. The postprocessor does not detect inconsistencies among files it operates on. For example, suppose you modify and recompile a class without also recompiling its subclasses. This can cause inconsistencies, which the postprocessor does not detect when it annotates the class files.

Modifications to Superclasses

When you run the postprocessor to make classes persistence capable, it generates annotated class files for the specified classes and for any superclasses that are in the same packages as the specified classes. ObjectStore requires annotations to superclasses for all classes that the postprocessor makes persistence capable. If a superclass is not in the same package as one of its subclasses that is being made persistence capable, you must explicitly specify the superclass on the postprocessor command line.

Effects on Inheritance

If a class that the postprocessor is annotating has no superclass other than `java.lang.Object`, the postprocessor annotates the class to implement the `com.odi.IPersistent` interface.

Implementation of the `IPersistent` interface is mandatory for objects that you want to be persistent. You must define classes so that if they inherit from another class, it is a class that can implement `IPersistent`.

Every class inherits from the `Object` class, which defines the `hashCode()` method and provides a default implementation. For a persistent object, this default implementation often returns a different value for the same persistent object (the object on the disk) at different times. This is because ObjectStore fetches the persistent object into different Java objects at different times, in different transactions or different invocations of Java.

This is not a problem if you never put the object into a persistent hash table or other structure that uses the `hashCode()` method to locate objects. If you do put them in hash tables or something similar, the hash table or other structure that relies on the `hashCode()` method might become corrupted when you bring the objects back from the database.

To resolve this problem, you can define your own `hashCode()` method and base it on the contents of the object so it returns the same thing every time. The signature of this method must be

```
public int hashCode()
```

If you do not provide a `hashCode()` method, the postprocessor adds one if it is necessary. If the default behavior of the postprocessor is not ideal for your application, you can specify the `-hashcode` and `-nohashcode` options to control where the postprocessor adds a `hashCode()` method.

Location of Annotated Class Files

When you run the postprocessor, you must specify a destination directory with the `-dest` option. The postprocessor uses the destination directory as the root directory of the class hierarchy of annotated files. The postprocessor places the annotated class file in the package-relative subdirectory of the destination directory. With the destination directory specified in your `CLASSPATH` environment variable, Java can find the annotated classes.

You must create the destination directory before you specify it in an `osjcfp` command line. The postprocessor creates the required subdirectories in the destination directory.

For example, suppose that you specify `osjcfpout` as the destination directory. When you run the postprocessor on the `Person.class` file, which is in the `com.odi.demo.people` package, the postprocessor places the annotated file in

```
osjcfpout\com\odi\demo\people\Person.class
```

The package name of the annotated class file remains the same, unless you specify an option to change it. The class name of the annotated class file is always the same as the class name of the unannotated class file.

Postprocessor Errors and Warnings

If an error occurs while the postprocessor is running, it terminates without writing any annotated class files.

For any warnings from the postprocessor, you might determine that you can safely ignore the warning. In this case, you can stop the postprocessor from warning you about the field in question. To do so, specify the `-quietfield` option followed by the fully qualified name of the field for which you want to suppress warnings. Alternatively, you can specify `-quietclass` to suppress all warnings on the class.

Handling of final Fields

You cannot make `final` fields persistent. If you try to do this, the postprocessor displays a warning message and treats the fields marked as `final` as though you declared them to be `transient`. To allow such fields to be stored persistently, you must remove the `final` keyword.

Handling of Static Fields

The postprocessor never stores static fields in the database and never causes the values of static fields to be altered. You must write your own code to update static fields and to store static fields in the database, if that is what you want to do.

Static fields that can hold persistent values

The postprocessor displays a warning for a static field that can hold potentially persistent values. The postprocessor cannot determine the type of the object that is actually pointed to. Consequently, depending on the type of object referenced, the warning might not be applicable. For example, suppose you have a persistence-capable class named `x`. The `x` class has a static member named `y` of a type that implements `com.odi.IPersistent`. When you run the postprocessor, it displays a warning such as the following:

```
x.y is a static field of a type that implements
com.odi.IPersistent and that might refer to a persistent object.
If this field does refer to a persistent object it must be user
maintained.
```

Referring to a persistent object

If the field mentioned in the warning is intended to refer to a persistent object, you can write your application as follows:

- When you create a database, create an object of the desired type and create a database root to refer to the object. This makes the object persistent and provides a mechanism that you can use to find the object.
- When you start a new transaction, if the object referenced by the static field is null or stale, look up the database root and set the value of the static field to the root value.

If you specify `ObjectStore.RETAIN_STALE` when you commit or abort a transaction, you must ensure that you correctly access the objects at the beginning of the next transaction. This is because `ObjectStore` does not make the object referenced by `x.y` persistent if it is only reachable from `x.y`. If `ObjectStore` makes it persistent because it is reachable from some other point, the object referenced by `x.y` might become stale at the end of the transaction in which it becomes persistent. If it does, and if the object referenced by `x.y` does become persistent, it is possible that the application might try to use the stale version of the object.

How can `x.y` be reachable from some other point? Perhaps another persistent object or an object that is going to be persistent refers to the object that the static data member is referring to. When `ObjectStore` commits the

transaction and performs transitive persistence, it finds the object that the static data member is referring to.

References to stale objects

An issue to consider is stale references to stale objects. To avoid the inadvertent use of stale objects, update `x.y` at transaction boundaries. Set `x.y` to null or to another value to ensure that if a stale object is referenced by `x.y`, it is no longer accessible through `x.y`. Then you can suppress the warning with the `-quietfield` option.

Summary

For class `x`, the important points are listed below.

```
class X {static OSHashtable y = new OSHashtable();}
```

- `x.y` does not become persistent just because class `x` is persistence capable or because an instance of `x` becomes persistent.
- If you want `x.y` to become persistent, you must make it reachable from a root through a path that does not involve a static field, for example, `db.createRoot("X.y", x.y)`.
- If `x.y` does become persistent, you must be aware that the `OSHashtable` object referenced by `x.y` might become stale at transaction boundaries. If it does, you must update `x.y` to refer to a nonstale instance.

Which Java Executable to Use

The postprocessor is a Java program; it requires a Java virtual machine to run. It uses the first Java executable that it finds in your `PATH` environment variable. If you want the postprocessor to use another Java executable, set the `OSJCFPJAV` environment variable to the name of the Java executable you want the postprocessor to use. The default is `java`.

If the postprocessor cannot find a Java executable, it generates a `Bad command or file name` error message.

Line-Number and Local-Variable Information

When the postprocessor annotates a class file, it maintains any existing line-number and local-variable information.

Using a Debugger

The class file postprocessor annotates methods with VM instructions for automatically performing `fetch()` and `dirty()` operations on objects. It does this in such a way that the debugging information in the class files remains intact. Typically, the annotations are invisible to an application. However, it is possible to encounter them under certain circumstances when

using a debugger. For example, you might encounter the following when you use the `Step into` command:

```
x = foo(y.m);
```

Stepping into that statement might cause you to enter the `ObjectStore` code that causes the contents of the `y` object to be fetched. In such a situation, use the `Step out` command to leave the `ObjectStore` code. Then use the `Step into` command again, which should then step into the call to the `foo()` method.

You should rely on the `Step over` command whenever possible. However, there are situations in which you must use the `Step into` command. If you inadvertently step into an `ObjectStore` method, step out of the `ObjectStore` code and return to your own code by doing one of the following:

- Use a `Step out` command.
- Set a breakpoint in the calling code.
- Use repeated `Step over` commands until the method returns.

Handling of `finalize()` Methods

The Java GC calls the `java.lang.Object.finalize()` method on an object that is no longer referenced. The GC does this before it frees the space occupied by the object. In this way, the `finalize()` method provides a hook that you can use to free resources that are not freed by garbage collection, for example, memory that was allocated by a native method call.

If your persistence-capable class defines a `finalize()` method (Technical Support recommends that it should not), the class file postprocessor inserts annotations at the beginning of the `finalize()` method that change the persistent object to a transient object. See [Avoiding `finalize\(\)` Methods](#) on page 148.

Description of Postprocessor Optimizations

The postprocessor optimizes `fetch()` and `dirty()` calls in several ways. If you determine that an optimization is preventing insertion of a required call to `fetch()` or `dirty()`, you can disable the optimization.

- For array objects in looping constructs, the postprocessor inserts the call to `fetch()` or `dirty()` only in the first loop iteration. To disable this optimization, specify the `-noarrayopt` option when you run the postprocessor. This causes the postprocessor to insert calls to `fetch()` or `dirty()` in every iteration.

- For constructors, the postprocessor does not insert full annotations that would allow constructors to handle modifications to newly constructed objects. To disable this optimization, specify the `-noinitializeropt` option when you run the postprocessor. This causes the postprocessor to fully annotate constructors so that they can correctly handle modifications to objects that become persistent during constructor execution.

If your application inserts objects into `ObjectStore` collections during construction of the objects being inserted, you must specify the `-noinitializeropt` option. Doing so avoids errors in the handling of modifications to the newly constructed objects.

- For access to fields relative to `this` in nonstatic member methods, the postprocessor optimizes calls to `fetch()` and `dirty()`. To disable this optimization, specify the `-nothisopt` option when you run the postprocessor. This causes the postprocessor to insert a `fetch()` or `dirty()` call for each access to a field in `this`.

You should disable these optimizations if you commit transactions or evict persistent objects as follows:

- If you call `commit()` or `evict()` while iterating over persistent array elements, specify `-noarrayopt` when you run the postprocessor.
- If you call `commit()` or `evict()` in between accesses to different fields of `this`, specify `-nothisopt` when you run the postprocessor.

If you want to disable all three optimizations, specify the `-noopt` option instead of the three individual options.

Including Transient and Already Annotated Classes

After you run the postprocessor, the annotated class files are in the package-relative subdirectory of the destination directory (root directory) you specified. You might want other class files in this destination directory. These could be transient (nonpersistence-capable or nonpersistence-aware) class files or files that have already been annotated.

Copying Classes to the Destination Directory

To copy certain files to the destination directory along with the annotated files, specify the `-copyclass` option followed by the name of the file you want to copy. For example:

```
osjcfp -dest osjcfpout a.jar -copyclass b.class
```

In this example, the postprocessor annotates the files in `a.jar` and copies them to the package-relative subdirectory of the `osjcfpout` directory. The postprocessor also copies `b.class` to the `osjcfpout` directory but it does not modify the `b.class` file.

You can follow the `-copyclass` option with one or more `.class` file names, class names, `.jar` file names, or `.zip` file names. This option applies to each name that follows it until the postprocessor reaches a `-pc` or `-pa` option.

Specifying Classes to Be Copied and Classes to Be Persistence Capable

Classes for which you specify the `-copyclass` option can overlap with classes for which you specify the `-persistcapable` or `-persistaware` option. For example:

```
osjcfp -dest osjcfpout -copyclass *.class -persistcapable a.class
```

This allows you to keep all files in a package together and annotate only the classes that need to be annotated. You need not partition classes into those that need annotations and those that do not. You can specify the same file with more than one persistence-mode option because the `-persistcapable` option and the `-persistaware` option override the `-copyclass` option.

When Can a Class Be Transient?

Suppose you have a persistence-capable class, class `A`. A class that refers to class `A` can be transient if all access to `A`'s nontransient data members is through methods on `A`. The methods of `A` will be properly annotated. Because all other classes only use `A`'s methods, the other classes do not need to be persistence-aware. Consequently, you need not postprocess any classes that refer to `A`.

Any class that directly accesses `A`'s nontransient data members must be either persistence capable or persistence aware. Any other class that refers to `A` and does not directly access nontransient data members can be transient. That is, you do not have to postprocess it.

An important exception to this is that if a class manipulates an array object that might be persistent (specifically, setting and getting array elements), that class must be annotated to be persistence aware. However, if the code that provides access to the array is annotated to access the values of the array, you can avoid making the class persistence aware. It is difficult to reliably implement this in the general case.

If you compile with optimization the classes that use the methods that get and set array values, the compiler might inline the get and set methods. In this case, you must make the class that uses the get and set methods persistence aware.

Putting Processed Classes in a New Package

Normally, the postprocessor places the annotated files in a package-relative subdirectory of the destination directory, and the annotated files have the same package names as the original files. However, there is an option that allows you to change the package name of files specified in the postprocessor command line. The `-translatepackage` option modifies the package name so that the persistence-capable version of the class is in one package and the transient version (the original) is in another package.

To help you use the `-translatepackage` option, this section discusses the following topics:

- Using the `-translatepackage` Option
- How the Postprocessor Applies the Option
- Updating References to New Package Name
- References to Transient and Persistent Versions of a Class
- References to Transient Instances of a Persistence-Capable Class

Using the `-translatepackage` Option

To create persistence-capable classes whose package name is different from the original package name, specify the `-translatepackage` option followed by the current package name, then the new package name. The format for this option is

```
{-translatepackage | -tp} orig_pkg_name new_pkg_name
```

For example, suppose you have the `a.b.C` class and you want to create the `d.e.C` persistence-capable class. Run the postprocessor in the following way:

```
osjcfp -dest osjcfpout -translatepackage a.b d.e C.class
```

Exact match
required

The specification for the original package name must exactly match the package name of the specified file. If there is not an exact match, the postprocessor does not place the annotated file in the new package. For example, suppose you have two classes named `com.odi.demo.New` and `com.odi.Old`. You want to move `com.odi.Old` to the `com.odi.beta` package and you specify the following command:

```
osjcfp -dest osjcfpout -tp com.odi com.odi.beta com.odi.demo.New com.odi.Old
```

The postprocessor places the annotated file for the `com.odi.Old` class in `com.odi.beta.Old` in the package-relative subdirectory of the `osjcfpout` directory (`osjcfpout\com\odi\betaa\com.odi.beta.Old.class`).

The postprocessor does not place the annotated file for `com.odi.demo.New` in a different package because the original package name is `com.odi.demo` and not just `com.odi`. The postprocessor annotates `com.odi.demo.New` and places it in `osjcfpout\com\odi\demo\com.odi.demo.New.class`.

How the Postprocessor Applies the Option

The postprocessor applies the `-translatepackage` specification to

- All classes in the original package that are locatable by means of the `CLASSPATH` environment variable or the `-classpath` option, if you specify it. The `-classpath` specification overrides the `CLASSPATH` environment variable.
- Files on the command line whose package name exactly matches the specification for the original package name. This is true for files processed with the `-persistcapable`, `-persistaware`, or `-copyclass` option.

When
copying files

It does not matter whether the postprocessor is making any other changes to the specified files. The postprocessor changes the package names of files for which the `-copyclass` option is specified along with new persistence-capable or persistence-aware files.

Multiple
option
specifications

You can specify this option more than once on a command line to specify several package translations. If you accidentally specify more than one translation for the same package, the postprocessor performs the last translation you specify in the command line.

Updating References to New Package Name

A change to the package name of a class requires updating all references to that class to reflect the new name.

The postprocessor updates the references in classes that it is currently operating on. This includes each class specified on the command line and each class found in a `.zip` file or `.jar` file that is specified on the command line.

The postprocessor cannot detect whether there are `.class` files for which the postprocessor was not called that refer to the renamed package. You must either run the postprocessor on the complete set of class files or modify the Java source of any files that the postprocessor is not annotating.

References to Transient and Persistent Versions of a Class

You might want a class to refer to both the transient and persistence-capable versions of another class.

It is not possible for the postprocessor to determine the references that should be to persistence-capable objects. Because of this, you must code the class so it uses the full path name of the different versions of the class. This is the only way to clarify the version of the class that is wanted. However, this technique works correctly only when you are operating across batches. It does not work when you are within the same batch.

Example

Following is an example of what that means. Suppose you have a utility class called `a.b.C`. You want to have both a transient and a persistence-capable version of `a.b.C`. When you run the postprocessor, you specify `-translatepackage` to create a persistence-capable version called `y.z.C`. Then in another class called `a.b.D`, you try to use both versions of the class. You write source code in `a.b.D` that explicitly refers to `y.z.C` much like the following:

```
int n= y.z.C.countThem()
```

When you try to compile `a.b.D`, compilation can succeed only if you put the annotated classes into the class path of the compiler. Otherwise, the compiler reports an error, because there is no such thing as `y.z.C`.

Also, it is not possible for `a.b.C` and `a.b.D` to be in the same batch, because the `-translatepackage` option would apply to `a.b.D`. This would make all of `a.b.D`'s calls go to the persistence-capable version, which is not what you want.

Steps to follow

To use persistence-capable and transient versions of the same class, follow these steps:

1 Create a utility library.

This is the first batch. This library creates transient versions of the class.

2 Run the postprocessor on the first batch and specify options that put the two different versions of the class in two different packages.

This step creates the persistence-capable version of the class.

3 Use the library from an application.

The application is the second batch.

4 Compile the application with the annotated files of the first batch, but not the second batch, in the compiler's class path.

References to Transient Instances of a Persistence-Capable Class

You can use instances of a persistence-capable class in a transient-only manner. No special action is required and the calls to `ObjectStore.fetch()` and `ObjectStore.dirty()` do nothing.

There is no need for the unannotated version of the class to be available at run time.

To use the annotated version of the class, even if you are using it transiently, the `osji.jar` or `stublib.jar` file must be available in the `CLASSPATH` at run time. If you are using the class only transiently, it can be the `stublib.jar` that is available.

Creating Persistence-Capable Classes with Transient Fields

You can create a persistence-capable class with transient fields. A transient field is a field that is not stored in the database. The postprocessor ignores transient fields. Use the `transient` keyword to create a transient field. For example:

```
class A {  
    transient java.awt.Component myVisualizationComponent;  
    int myValue;  
}
```



```
...
}
```

In this class, the `myVisualizationComponent` field is declared to be a transient reference to `java.awt.Component`. The `java.awt` package contains GUI classes that do not lend themselves to being persistence capable.

In your persistence-capable class, you might have transient fields that you want to be able to access outside a transaction. In this situation, you can specify the `-noannotatefield` or `-naf` option for the field when you run the postprocessor. This option prevents access to the specified field from causing `fetch()` and `dirty()` calls on the containing object. Normally, access to a transient field causes `fetch()` or `dirty()` to be called to allow the `postInitializeContents()` and `preFlushContents()` methods to convert between persistent and transient state.

Transient Fields and Serialization

If you have a class that has fields that are declared as transient, this causes the default handling of these fields by object serialization to ignore the fields. If you want them ignored by object serialization and you want them to be stored persistently, specify the `-ignoretransient` option for the class when you run the postprocessor.

On the other hand, there might be a field that must be available for object serialization, but you do not want to store that field in the database. In this situation, specify the `-transientfield` option for the field when you run the postprocessor. This option causes the postprocessor to treat the specified field as though it has a `transient` modifier, even if it does not.

Initialization of Some Transient Fields

In the declaration of a transient field in a persistence-capable class, you might want to initialize the value of the transient field. However, when the postprocessor creates the hollow object constructor for the class, it does not define the constructor to initialize the transient field. This is true even when you specify the `final` keyword. The postprocessor does not initialize such fields because the initialization occurs as inlined code in each of the constructors for the class. For example:

```
private transient final MyField myField = new MyField();
```

The `final` keyword indicates to the postprocessor that initialization is required. However, the initialization code is not readily available and `myField` is not initialized. There are several ways to handle this situation.

You can create the hollow object constructor manually. For example, suppose you define the `MyField` class, which extends the `MyFarm` class, as in the following example:

```
...
public MyField(com.odi.ClassInfo dummyClassInfo) {
    super(dummyClassInfo);
}
```

This requires you also to manually define a hollow object constructor for the `MyFarm` class and for each superclass of the `myFarm` class.

Alternatively, you can remove the `final` qualifier and initialize the transient field in an `IPersistent.postInitializeContents()` method.

If you include an inline initialization of a field declared to be `transient` and `final`, the postprocessor displays an error message and stops processing. If you include an inline initialization of a field declared to be `transient`, but not `final`, the postprocessor warns you about the situation and continues processing. If you determine that you can safely ignore the message, you can turn it off with the `-ignoretransient` option to the postprocessor.

See also *Transient Fields in Persistence-Capable Classes* on page 146.

Customizing Updated Classes

There are several ways you can customize persistence-capable and persistence-aware annotations: You can implement your own versions of methods that the postprocessor typically adds; you can implement hook methods that `ObjectStore` calls at specified points; you can define a hollow object constructor in place of the hollow object constructor the postprocessor typically defines; you can also insert your own `fetch()` and `dirty()` calls.

Implementing Customized Methods and Hook Methods

The three methods described next are among the several annotations that the postprocessor adds to persistence-capable classes.

- The `initializeContents()` method loads real values into hollow instances of your persistence-capable class. In other words, hollow objects become active objects with an internal clean state.

- The `flushContents()` method copies values from a modified instance (active persistent object) back to the database. This changes the internal clean or dirty state of the persistent object to the clean state.
- The `clearContents()` method resets the values of an instance to the default values. This changes a clean active object to a hollow object.

Alternatives

If you want to, you can customize the behavior of these methods in the following two ways:

- Implement the method yourself. See Defining Required Methods in the Class Definition on page 239. If you do, the postprocessor does not add the method. However, if you implement any of the three methods listed previously, you must implement all of them. Also, you must define the `ClassInfo` subclass, define an instance of it, and register the instance. This is because the `ClassInfo` instance and the three previous methods must agree on the conventions for field numbering. An example of a program that implements these methods is in the `com/odi/demo/rep` directory in the `Rectangle.java` file. See `com/odi/demo/rep/README.htm`.
- Implement the hook method that corresponds to the method you want to customize. The postprocessor does not annotate hook methods. These hook methods provide a way to perform transient field maintenance. You might also be able to use these methods as an update mechanism for notification about a change:
 - `postInitializeContents()` — If you define this method, `ObjectStore` calls it immediately after it calls the `initializeContents()` method.
 - `preClearContents()` — If you define this method, `ObjectStore` calls it just before it calls the `clearContents()` method.
 - `preFlushContents()` — If you define this method, `ObjectStore` calls it just before it calls the `flushContents()` method.

Warning

The body of a hook method must not call any methods of the class and must not start or end a transaction. This is because the class methods are annotated and, therefore, make calls to `fetch()` and `dirty()`. Such calls in the middle of initializing or writing the object are not allowed because they might cause the virtual machine to encounter a stack overflow.

Sample program with hook methods

Following is an example of a program that implements these hook methods:

```
import com.odi.*;

/**
 * PColor provides a persistent representation of colors that can
 * be used with the Java AWT package. The java.awt.Color class
 * itself cannot be stored persistently, because some of its
```

```

* internal state depends on the particular kind of color display
* being used. If a java.awt.Color were created on a computer
* that used a 24-bit-deep color monitor, stored in a database,
* and then retrieved and used on a different computer that had
* a gray-scale monitor, it would not function correctly. PColor
* stores the color value as three integers, and then recreates
* the java.awt.Color object whenever the PColor object is
* brought into Java from persistent storage.
*
* For expository purposes, this example pretends that the value
* of a java.awt.Color object can change after the object is
* created. The real java.awt.Color class is immutable, and so
* the setBlue method below would not work, and the
* preFlushContents method would not actually be needed.
*/

public class PColor {

    /*These instance variables are stored persistently. They
    represent the color value. */
    int red;
    int green;
    int blue;

    /*This instance variable is declared transient, so it is not
    stored persistently. It is managed by the methods below. */
    transient java.awt.Color color;

    PColor(int r, int g, int b) {
        red = r;
        green = g;
        blue = b;
        color = new java.awt.Color(r, g, b);
    }

    /*When a PColor is brought into Java from persistent storage,
    the java.awt.Color object is created. Note that this method
    runs after the initializeContents, so that it can use the
    values of the persistent instance variables. */

    public void postInitializeContents() {
        color = new java.awt.Color(red, green, blue);
    }

    /*When a PColor is sent out from Java to persistent storage, the
    color value from the java.awt.Color object is copied into the
    persistent instance variables, so that it will be saved.
    Note that this method runs before flushContents, so that it
    can set up the values of the persistent instance variables. */

    public void preFlushContents() {
        red = color.getRed();
        green = color.getGreen();
        blue = color.getBlue();
    }

    /*When clearContents happens, this method sets the color

```

```

instance variable to null, so that this PColor object won't be
stopping the java.awt.Color object from being reclaimed. */

    public void preClearContents() {
        color = null;
    }
    /*Equality for PColor objects is the same as equality of the
    underlying java.awt.Color objects. */

    public boolean equals(Object obj) {
        if (obj instanceof PColor) {
            return color.getRGB() == ((PColor)obj).color.getRGB();
        }
        return false;
    }
    public java.awt.Color getColor() {
        return color;
    }
    public int getBlue() {
        return color.getBlue();
    }
    public int setBlue(int b) {
        color.setBlue(b);
    }
    /* and so on.... */
}

```

Creating a Hollow Object Constructor

For each persistence-capable class, the postprocessor finds or generates a hollow object constructor. The hollow object constructor takes a single argument whose type is `com.odi.ClassInfo`. Typically, you need not define a hollow object constructor, but you can if you want to.

Why define one?

A reason to define your own hollow object constructor is to initialize transient fields that you want to be usable, even if the `fetch()` method has not been called.

You should avoid performing actions in a hollow object constructor that would cause the object to be fetched. Doing so might cause infinite recursion to occur.

For example, if a class has a persistent `hashCode()` method, it is a bad idea to define a hollow object constructor to register the instances of the class in a hash table. Doing so would cause the `hashCode()` method to be called, which in turn would attempt to fetch the object.

Creation steps

When the postprocessor creates the hollow object constructor, it follows these steps:

- 1 The postprocessor selects an appropriate superclass hollow object constructor.

If the superclass has an accessible constructor that takes a single `com.odi.ClassInfo` argument, or if it will have one because the postprocessor adds it during this execution of the tool, the postprocessor uses that constructor. The postprocessor reports an error if it cannot find an accessible constructor.

- 2 The postprocessor creates a public constructor that
 - Accepts a `com.odi.ClassInfo` argument
 - Invokes the selected superclass constructor
 - Initializes all persistent fields to an appropriate default state that is equivalent to the result of the `clearContents()` method

You can define the hollow object constructor instead of allowing the postprocessor to do it. If you define one, the postprocessor does not generate one.

Optimizing Operations That Retrieve Persistent Objects

Before an application can access the contents of a persistent object, it must call the `ObjectStore.fetch()` method to read the object or the `ObjectStore.dirty()` method to modify the object. These calls make the contents of the object available to your application. The postprocessor inserts these calls in methods of classes that it makes persistence capable or persistence aware. However, the postprocessor might not annotate your code for best performance. You might find that you can improve performance by inserting the `fetch()` and `dirty()` calls yourself.

Caution If you insert a `fetch()` or `dirty()` call in a method, the postprocessor does not add any additional `fetch()` or `dirty()` calls to that method.

Procedure for Optimizing Operations

Before you add the calls yourself, first allow the postprocessor to add the `fetch()` and `dirty()` calls. Then run and monitor your program. If you want to try to improve performance, add the calls to your source file and recompile. When you run the postprocessor again, it recognizes that the

`fetch()` or `dirty()` call is already in place and does not add any `fetch()` or `dirty()` calls to any methods that already contain such a call.

If you do this annotation, you should also add `implements IPersistent` to the definition of any class that is accessed with a `fetch()` or `dirty()` call. When you do this, the compiler can effectively use the multiple overloads of the `fetch()` and `dirty()` methods, which take `com.odi.IPersistent` arguments. Also, the compiler can generate more efficient code when you declare the class to implement `IPersistent` in your source.

Inlining Code

An important consideration when annotating by hand is that the compiler might inline the code into calling methods. This makes it appear to the postprocessor that the code annotations are in the calling method, which might not be true.

When you are using the JDK `javac` compiler, this occurs when you specify the `-O` (capital O, as in Oslo) option.

To ensure that the postprocessor functions correctly, you must do one of the following:

- Prevent the compiler from inlining code.
- If you add `fetch()` and `dirty()` calls to a method that is a candidate for inlining, also annotate all the methods that call that method. A method is a candidate for inlining if it calls `static`, `final`, or `private` methods, or invokes methods with the `super` qualification construct.

Preventing Fetch of Transient Fields

You might want to avoid the insertion of the `fetch()` call in methods that operate only on transient fields. A strategy for doing this takes advantage of the fact that the postprocessor does not annotate a method if it already includes a `fetch()` or `dirty()` call. If you know that a method operates only on transient fields, you can prevent insertion of the `fetch()` call with code such as the following:

```
try {
    method body goes here
} catch (SomeRuntimeExceptionThatWillNotOccur) {
    ObjectStore.fetch(this);
}
```

This imposes no execution time and prevents the postprocessor from inserting the `fetch()` method. You can create your own exception, which

inherits from `java.lang.RuntimeException`, or select an existing one. The safest approach is to create your own exception so you can be sure that the exception is never signaled.

Performing a Test Run of the Postprocessor

You can run the postprocessor without actually updating any files. The tool performs all processing and error checking and can display messages that indicate what it is doing. This allows you to make corrections before creating the persistence-capable versions of your classes.

To perform a test run of the postprocessor, specify the `-nowrite` option on the command line. For example:

```
osjcfp -dest osjcfpout -nowrite classes.jar
```

This command processes all class files in the `.jar` file and displays any error messages. To view information messages from the postprocessor, include the `-verbose` option. For example:

```
osjcfp -dest osjcfpout -nowrite -verbose classes.jar
```

It does not matter where you place the `-nowrite` or `-verbose` option in the command line. Wherever you place them, they apply to all files that the postprocessor processes.

To suppress nonfatal warning messages, specify the `-quiet` option. The `-quiet` and `-verbose` options are mutually exclusive. The last one used on the command line applies to the entire execution. For example, the following line suppresses warning messages during the processing of all specified files because the `-quiet` option follows the `-verbose` option.

```
osjcfp -dest osjcfpout -nowrite -verbose classes.jar -quiet more.jar
```

You can also suppress some, but not all, warnings. Specify the `-quietclass` option followed by the fully qualified name of a class to suppress warnings for that class. Specify the `-quietfield` option followed by the fully qualified name of a field to suppress warnings that pertain to that field. These options apply only to the element whose name immediately follows the option. If the `-verbose` option is also specified, these options take precedence.

Using an Input File

When you are running the postprocessor on a lot of files and specifying many options, the command line can be very long. As a convenience, you can enter the options and file names in a file, then specify the file name as a postprocessor option. Be sure to prefix the file name with the @ symbol.

Windows

On Windows systems, there is a limit of eight arguments on a command line. Consequently, you usually must use input files on Windows.

Format

You can include comments in the input file. You can place items on different lines and line continuation symbols are not required. Line breaks are treated as white space. Otherwise, enter data in the input file exactly as you would enter it on the command line.

Indicate comments with a # sign. The postprocessor ignores any subsequent characters on the same line as the # sign.

Example

For example, suppose you enter some postprocessor options and files for the postprocessor to operate on in an input file named `optionsAndFiles`. You specify this file as follows:

```
osjcfp @optionsAndFiles
```

You can intersperse input file specifications with options and files that you enter on the command line. For each specified input file, the postprocessor removes any comments from the input file and replaces the input file specification with the data in the input file. The postprocessor then begins to process the command line. For example:

```
osjcfp -dest osjcfpout @file1 -tp old.pack new.pack @file2
```

The postprocessor

- 1 Replaces `@file1` with the contents of `file1`
- 2 Replaces `@file2` with the contents of `file2`
- 3 Executes the command line starting with the `-dest` option

Nesting and wildcards

You cannot nest input file specifications. That is, you cannot include the `@file_name` option in an input file. Also, you cannot use wildcards in an input file. The postprocessor does not expand them.

Annotations You Must Add

There are some annotations that the postprocessor either cannot perform or does not perform because of execution performance considerations. You must include these annotations when you code your source files.

Keep in mind that when you add even one `fetch()` or `dirty()` call to a method, the postprocessor recognizes that the method is already annotated and does not add any other `fetch()` or `dirty()` calls to that method. If you do annotate a method, be sure to add all required calls.

This section provides information about the following topics:

- Interfacing with Nonpersistent Methods
- Interfacing with Native Classes
- Annotating Subclasses
- Passing Arrays
- Implementing the Hollow Object Constructor for Some Instance Fields
- Using the Java Reflection API with Persistence-Capable Objects

Interfacing with Nonpersistent Methods

It is possible for a method in a persistence-capable class to pass a persistent object to a nonpersistent method. When this happens, you must ensure that there is a `fetch()` or `dirty()` call for the persistent object before it is passed to the nonpersistent method.

If all access to persistent objects is through annotated methods (methods in persistence-capable or persistence-aware classes), manual annotations are not required. For arrays, there is no way to define a class so arrays of that class can be accessed only by persistence-aware classes. You must be sure to call the `fetch()` or `dirty()` method on a persistent array before passing it to a method in a nonpersistent class.

Interfacing with Native Classes

The postprocessor cannot analyze or annotate native methods. If your code passes a persistent object to a native method, and if the native code might try to access the object other than through annotated methods, be sure to insert a call to `fetch()` or `dirty()` for the persistent object before it is passed. In cases in which native code might access or navigate among persistent objects, you must do one of the following:

- Modify the native code to call `fetch()` or `dirty()` itself.
- Make the necessary `fetch()` and `dirty()` calls before calling the native method.

Annotating Subclasses

After you create a persistence-capable or persistence-aware class, you can define a subclass of that class. Doing so does not make the subclass persistence capable or persistence aware. You must run the postprocessor on the subclass.

If you forget to run the postprocessor on a subclass and if the subclass is reachable from a persistent root, other than through a transient field, `ObjectStore` might try to migrate instances of the subclass to the database. This attempt causes an error because the subclass is not persistence capable.

Passing Arrays

In your application, you might pass an array to a nonpersistent method when the nonpersistent method is defined as having a parameter of type `java.lang.Object`. In this situation, the postprocessor cannot determine that it should insert `fetch()` or `dirty()` calls for the array in the calling method before passing the array. You must annotate the calling method yourself.

If the called method is declared to accept an array argument, the postprocessor recognizes that a `fetch()` call might be needed and inserts it.

Implementing the Hollow Object Constructor for Some Instance Fields

A class can include nonstatic (instance) fields that contain initializer expressions in their declarations. Postprocessor-generated `ClassInfo` constructors do not run these initializers. Normally, this is not a problem. The constructor allows hollow object initialization and the `initializeContents()` method overwrites these fields when the object is fetched.

However, there might be transient nonstatic fields that have initializer expressions or fields that are treated as transient by your implementation of the `ClassInfo` type and the `initializeContents()` and `flushContents()` methods. In this case, you must manually implement the hollow object constructor or `ObjectStore` does not run the initializer. It is impossible for the

postprocessor to detect such cases, and no warning message can be provided. See [Creating a Hollow Object Constructor](#) on page 229.

Using the Java Reflection API with Persistence-Capable Objects

You can use the `java.lang.reflect.Field` class to get and set fields of persistence-capable objects. To do so, you must

- Call the `ObjectStore.fetch()` method for an object before you call any of the `java.lang.reflect.Field` *get* methods to get the value of any of the object's fields.
- Call the `ObjectStore.dirty()` method for an object before you call any of the `java.lang.reflect.Field` *set* methods to set the value of any of the object's fields.

Class File Postprocessor Limitations

It is possible to cause invalid references when you run the postprocessor and rename the package. In an annotated class, the postprocessor locates and updates class names if they are in field, method, or class references. The postprocessor cannot locate and update string arguments to `Class.forName()` if the name specifies a class whose package has been renamed.

Chapter 9

Generating Persistence-Capable Classes Manually

This chapter provides information about the way to define persistence-capable and persistence-aware classes in your program explicitly without using the automated class file postprocessor supplied with ObjectStore.

Technical Support recommends that you use the automated postprocessor. See Chapter 8, *Generating Persistence-Capable Classes Automatically*, on page 195. For information on which Java-supplied classes are persistence capable, see *Java-Supplied Persistence-Capable Classes* on page 311.

You might choose the manual method if you want to

- Manually optimize the code
- Perform translation between nonpersistent objects and a custom persistent representation

You can partially manually annotate a class, then run the postprocessor to insert the remaining required annotations.

You must explicitly postprocess or manually annotate each class that you want to be persistence capable. The capacity for an object to be stored in a databases is not inherited when you subclass a persistence-capable class.

Contents

This chapter discusses the following topics:

Explicitly Defining Persistence-Capable Classes	238
Additional Information About Manual Annotation	247
Creating and Accessing Fields in Annotations	253

Explicitly Defining Persistence-Capable Classes

Follow these steps to annotate your program so that classes you define are persistence capable:

- 1 Define your class to implement the `IPersistent` interface. See page 238.
- 2 In the class definition, define the required fields. See page 239.
- 3 In the class definition, define the required methods. See page 239.
- 4 In the class definition, define accessor methods so that they make the appropriate `ObjectStore.fetch()` and `ObjectStore.dirty()` method calls. See page 242.
- 5 If required, define a class that extends the `ClassInfo` class. See page 243.

Interfaces never require `ClassInfo` classes.

If you will be running your application in an environment that allows the unrestricted use of the Java reflection API, public or abstract classes with hollow object constructors do not require `ClassInfo` classes. However, all classes that define indexable fields on objects stored in peer (`com.odi.coll`) collections do require `ClassInfo` classes.

- 6 For any `ClassInfo` subclasses you define, create an instance of the `ClassInfo` subclass. Only one instance of this subclass is ever needed. See page 244.
- 7 Call the static `get()` method on `ClassInfo`. (Typically, this is in static initializer code for the manually annotated class.) See page 244.

Some Java-supplied classes are persistence capable. Others are not persistence capable and cannot be made persistence capable. A third category of classes can be made persistence capable, but there are important issues to consider when you do so. Be sure to read *Java-Supplied Persistence-Capable Classes* on page 311.

About interfaces

Interfaces are always persistence capable. You must specify them when you run the postprocessor, but other than that, you need not do anything to make an interface persistence capable.

Implementing the IPersistent Interface

Every persistence-capable class must implement the `IPersistent` interface or be a subclass of a class that implements it. As with any interface, every

method defined in the `IPersistent` interface must be defined in a class that implements `IPersistent`. If you want to rely on the postprocessor to insert the missing methods for you, you must not explicitly implement the `IPersistent` interface.

Defining the Required Fields

The following code must be in your class definition. You can add this code yourself, or you can run the postprocessor to add it.

```
transient private com.odi.imp.ObjectReference ODIDef;
transient public byte ODIObjState;
```

The `ODIDef` field stores a reference. The `ODIObjState` field holds some object state bits. The underlying run-time classes in `ObjectStore` access these fields through the `IPersistent` accessor methods, as needed.

Defining Required Methods in the Class Definition

This section describes the methods that must be defined in a class that implements the `IPersistent` interface.

Define the `initializeContents()` method to load real values into hollow instances of your class. This changes a hollow object to an active object. `ObjectStore` provides methods on the `GenericObject` class that retrieve each `Field` type. Be sure to call the correct methods for the fields in your persistent object. There is a separate method for obtaining each type of `Field` object. `ObjectStore` calls the `initializeContents()` method as needed. The method signature is

```
public void initializeContents(GenericObject genObj)
```

Following is an example:

```
public void initializeContents(GenericObject handle) {
    name = handle.getStringField(1, PCI);
    age = handle.getIntField(2, PCI);
    children = (Person[])handle.getField(3, PCI);
}
```

If the class you are annotating implements `IPersistent` through a superclass, you must also initialize superclass fields by invoking `initializeContents()` on the superclass.

Define the `flushContents()` method to copy values from a modified instance (active persistent object) back to the database. This method changes an active clean or dirty object to an active clean object. `ObjectStore` provides methods on the `GenericObject` class that set each `Field` type. Be sure to call

the correct methods for the fields in your persistent object. There is a separate method for setting each type of `Field` object. `ObjectStore` calls the `flushContents()` method as needed. The method signature is

```
public void flushContents(GenericObject genObj)
```

Following is an example:

```
public void flushContents(GenericObject handle) {
    handle.setClassField(1, name, PCI);
    handle.setIntField(2, age, PCI);
    handle.setArrayField(3, children, PCI);
}
```

If the class you are annotating implements `IPersistent` through a superclass, you must also flush superclass fields by invoking `flushContents()` on the superclass.

Define the `clearContents()` method to reset the values of an instance to the default values. This method changes an active clean object to a hollow object. This method must set all reference fields that referred to persistent objects to null. `ObjectStore` calls this method as needed. The method signature is

```
public void clearContents()
```

Following is an example:

```
public void clearContents() {
    name = null;
    age = 0;
    children = null;
}
```

If the class you are annotating implements `IPersistent` through a superclass, you must also clear superclass fields by invoking `clearContents()` on the superclass.

Field accessor methods

The following accessor methods must be in the class definition:

- `public ObjectReference ODIgetRef()`
- `public void ODIsetRef(ObjectReference objRef)`
- `public byte ODIgetState()`
- `public void ODIsetState(byte state)`

If you do not want to define them, you can run the postprocessor to insert them for you, but you must not declare the class to implement `IPersistent`. However, if you explicitly define an `ODIgetxxx()` method, you must explicitly define its associated `ODIsetxxx()` method. Likewise, if

you explicitly define an `ODIsetxxx()` method, you must explicitly define its associated `ODIgetxxx()` method.

If you add the code yourself, it must look like the following:

```
public com.odi.imp.ObjectReference ODIgetRef() {
    return ODIRef;
}

public void ODIsetRef(com.odi.imp.ObjectReference objRef) {
    ODIRef = objRef;
}

public byte ODIgetState() {
    return ODIObjState;
}

public void ODIsetState(byte state) {
    ODIObjState = state;
}
```

Implementing the IPersistentHooks Interface

There are times when you might want a persistence-capable class to maintain transient information in parallel with persistent information. The `IPersistentHooks` interface allows you to do this.

If a persistence-capable class implements the `IPersistentHooks` interface, `ObjectStore` calls the `IPersistentHooks` methods that are defined when it calls the corresponding methods defined in the `IPersistent` interface.

As with any interface, every method defined in the `IPersistentHooks` interface must also be defined in the class that explicitly implements the `IPersistentHooks` interface.

If you explicitly declare that a class implements the `IPersistentHooks` interface without providing all the definitions declared in the interface, you receive a compilation error. However, if you define some methods, but not all of them, and you do not explicitly declare that the class implements the `IPersistentHooks` interface, you can use the postprocessor to insert the missing methods and the interface declaration.

Hook methods

The following methods must also be in the class definition. You can define them as methods with empty bodies. If you do not define them and your class does not explicitly implement the `IPersistentHooks` interface, you can use the postprocessor to add these methods with empty bodies.

- `postInitializeContents()` is called by `ObjectStore` immediately after it calls the `initializeContents()` method.

- `preFlushContents()` is called by `ObjectStore` immediately before it calls the `flushContents()` method.
- `preClearContents()` is called by `ObjectStore` immediately before it calls the `clearContents()` method.
- `preDestroyPersistent()` is called by `ObjectStore` immediately before it calls the `ObjectStore.destroy()` method.

Making Object Contents Accessible

In each class that you want to be persistence capable, you must annotate your class definition to include calls to the `ObjectStore.fetch()` and `ObjectStore.dirty()` methods. It does not matter whether the class explicitly implements `IPersistent` or inherits from a class that implements `IPersistent`. These calls are required for the class to be persistence capable.

With some exceptions, before your application can access the contents of an object, it must call the

- `ObjectStore.fetch()` method on the object to read its contents
- `ObjectStore.dirty()` method on the object to modify its contents

Calls to
`fetch()` or
`dirty()`

Your application calls the method and passes an object whose contents you want to access. This makes the contents of the object available. Modify the methods that reference nonstatic fields to call the `ObjectStore.fetch()` and `ObjectStore.dirty()` methods as needed. While this step is not mandatory, it does provide a systematic way of ensuring that the application calls the `fetch()` or `dirty()` method before accessing or updating object contents.

Remember that you can add some annotations and run the postprocessor to add other annotations. You might want to define the required methods and the `ClassInfo` subclass, but let the postprocessor insert the required `fetch()` and `dirty()` calls.

Exceptions

You need not call the `fetch()` or `dirty()` method on instances of primitive wrapper classes (see *Description of Java-Supplied Persistence-Capable Classes*). If you do call `fetch()` or `dirty()` on these objects, nothing happens and processing continues.

You need not call the `fetch()` or `dirty()` method on instances of `java.lang.String`. You need not call the `fetch()` or `dirty()` method on the following objects. If you do call `fetch()` or `dirty()` on these objects, nothing happens and processing continues.

- Instances of primitive wrapper classes
- Java peer objects (remember that `ObjectStore` collection objects are Java peer objects)

If you call `fetch()` on instances of `java.lang.String`, nothing happens. If you call `dirty()` on instances of `java.lang.String`, `ObjectStore` signals `ObjectException`.

Defining a `ClassInfo` Subclass

If required, define a public class that inherits from the `ClassInfo` class. (See page 247 for requirements.) You must define this class in a separate file. If you plan to use the postprocessor to insert any annotations, the name of this class must be one of the following:

- The name of the persistence-capable class, followed by `ClassInfo`, for example, `PersonClassInfo`
- The suffix specified with the `-classinfosuffix` option to the postprocessor

In each `ClassInfo` subclass definition, you must include the following methods.

Define a `create()` method to create instances of your persistence-capable class with default field values:

```
public IPersistent create() {return new Person(this);}
```

This should call a constructor, referred to as a hollow object constructor, that leaves fields in the default state. For an abstract class, the `create()` method can return null.

Define the public `getClassDescriptor()` method to obtain the class object for your class. For example:

```
public Class getClassDescriptor()
    throws ClassNotFoundException {
    return Class.forName("com.odi.demo.people.Person"); }
```

Define the public `getFields()` method to allow access to the names and types of the fields of the class. For example:

```
public Field[] getFields() { return fields; }
private static Field[] fields = {
    Field.createString("name"),
    Field.createInt("age"),
    Field.createClassArray("children", "Person", 1)
};
```

The definition of the `getFields()` method can specify create methods for fields that are not in the class definition and can omit create methods for fields that are in the class definition.

Example of a Manually Annotated Persistence-Capable Class

Following is an example of a definition of a manually annotated persistence-capable class. Three consecutive periods indicate lines from a complete program that have been omitted here because they are not pertinent to creating a persistence-capable class.

**Class
definition**

```
package com.odi.demo.people;
import com.odi.*;

// Define a class that implements IPersistent:
class Person implements IPersistent {

    // Fields:
    String name;
    int age;
    Person children[];
    // Other fields ...

    // Constructor:
    public Person(String name, int age, Person children[]) {
        this.name = name; this.age = age; this.children = children;
    }

    // Hollow object constructor:
    public Person(ClassInfo info) { }

    // Accessor methods that have been modified to call
    // the fetch() and dirty() methods:
    public String getName() {ObjectStore.fetch(this);
        return name; }
    public void setName(String name) {ObjectStore.dirty(this);
        this.name = name; }
    public int getAge() {ObjectStore.fetch(this); return age; }
    public void setAge(int age) {ObjectStore.dirty(this);
        this.age = age; }
    public Person[] getChildren() {ObjectStore.fetch(this);
        return children; }
    public void setChildren(Person children[]) {
        ObjectStore.dirty(this); this.children = children;
    }
    // Other methods ...

    // Additions required for ObjectStore:
```

```

// Define the initializeContents() method to load real
// values into hollow persistent objects, which makes
// them active persistent objects:

public void initializeContents(GenericObject handle) {
    name = handle.getStringField(1, myClassInfo);
    age = handle.getIntField(2, myClassInfo);
    children = (Person[])handle.getField(3, myClassInfo);
}

// Define the flushContents() method to copy the
// contents of a persistent object to the database:

public void flushContents(GenericObject handle) {
    handle.setClassField(1, name, myClassInfo);
    handle.setIntField(2, age, myClassInfo);
    handle.setArrayField(3, children, myClassInfo);
}

// Define the clearContents() method to reset the values
// of a persistent instance to the default values.
// This method must set all reference fields that
// referred to persistent objects to null:

public void clearContents() {
    name = null;
    age = 0;
    children = null;
}

// Define the ODIRef and ODIObjState fields and
// their accessor methods.
transient private com.odi.imp.ObjectReference ODIRef;
transient public byte ODIObjState;
public com.odi.imp.ObjectReference ODIGetRef() {
    return ODIRef;
}
public void ODISetRef(com.odi.imp.ObjectReference objRef) {
    ODIRef = objRef;
}
public byte ODIGetState() {
    return ODIObjState;
}
public void ODISetState(byte state) {
    ODIObjState = state;
}

// Create an instance of the subclass of ClassInfo and
// register that instance:

static ClassInfo myClassInfo =
    ClassInfo.get("com.odi.people.Person");
}

```

**ClassInfo
definition**

In a separate file, define the subclass of the `ClassInfo` class if its definition is required. For example:

```
// Define the subclass of ClassInfo. A recommended naming
// convention is to prefix the name of your persistence-capable
// class to "ClassInfo".

package com.odi.demo.people;

import com.odi.*;

public class PersonClassInfo extends ClassInfo {

    // Define a create() method to create instances of your
    // class with default field values. The method
    // calls the hollow object constructor and passes this,
    // which is an instance of the ClassInfo subclass:

    public IPersistent create() { return new Person(this); }

    // Define these public methods to provide access to
    // the name of the persistence-capable class, the name of its
    // superclass, and the names of its fields.
    // The array returned by getFields() must contain the
    // fields in the order of their field numbers.

    public Class getClassDescriptor()
        throws ClassNotFoundException {
    return Class.forName("com.odi.demo.people.Person"); }

    public Field[] getFields() { return fields; }
    private static Field[] fields = {
        Field.createString("name"),
        Field.createInt("age"),
        Field.createClassArray(
            "children", "com.odi.demo.People.Person", 1)
    };
}
```

It does not matter whether the `ClassInfo` class explicitly implements `IPersistent` or inherits from a class that implements `IPersistent`.

`ClassInfo` is an abstract class for managing schema information for persistence-capable classes. `ObjectStore` requires the schema information to manage the object. If you do not explicitly define a `ClassInfo` class, `ObjectStore` uses the Java reflection API to create the needed information at run time.

After you perform the steps described in this section, you can store instances of your class in a database.

ObjectStore does not let you store `final` instance variables persistently. This is because it is not possible to write the `initializeContents()` and `clearContents()` methods to handle `final` instance variables correctly.

Additional Information About Manual Annotation

This section provides additional information about manually annotating a class to be persistence capable. It discusses the following topics:

- Defining a `hashCode()` Method
- Defining a `clone()` Method
- Working with Transient-Only and Persistent-Only Fields
- Defining Persistence-Aware Classes
- Following Postprocessor Conventions
- Annotating Abstract Classes

Defining a `hashCode()` Method

Every class inherits from the `Object` class, which defines the `hashCode()` method and provides a default implementation. For a persistent object, this default implementation often returns a different value for the same persistent object (the object on the disk) at different times. This is because ObjectStore fetches the persistent object into different Java objects at different times (in different transactions or different invocations of Java).

This is not a problem if you never use the object as a key in a persistent hash table or other structure that uses the `hashCode()` method to locate objects. If you do use the object as a key, the hash table or other structure that relies on the `hashCode()` method might become corrupted when you bring the objects back from the database.

To resolve this problem, you can define your own `hashCode()` method and base it on the contents of the objects so it returns the same thing every time. The signature of this method must be

```
public int hashCode()
```

Defining a clone() Method

If your persistence-capable class implements the `Cloneable` interface, your class must define a `clone()` method. This `clone()` method must ensure that it correctly initializes and checks the `ODIRef` and `ODIObjectState` fields when it performs a clone operation. For new cloned objects, your application should initialize `ODIRef` to null and `ODIObjectState` to zero.

Working with Transient-Only and Persistent-Only Fields

The definition of the `ClassInfo.getFields()` method returns an array of `com.odi.Field` instances. There is one element for each field that you want to store and retrieve in a persistent object. `ObjectStore` does not require an exact match between each field in the Java class definition and each field array element returned by the `getFields()` method. Furthermore, fields listed in the `getFields()` return value need not directly represent fields in the class. They can represent state from which values for fields in the class are synthesized.

Transient-only fields

A persistence-capable Java class can define a field that does not appear in the list of fields returned by the `ClassInfo.getFields()` method. Such a field is a transient-only field. The `initializeContents()` method that is associated with the class can be used to initialize transient-only fields based on persistent state. For example:

```
class A {
    transient java.awt.Component myVisualizationComponent;
    int myValue;
    ...
}
```

In this class, the `myVisualizationComponent` field is declared to be a transient reference to `java.awt.Component`. The `java.awt` package contains GUI classes that do not lend themselves to being persistence capable.

Number of fields

The number of nonstatic, nontransient declared fields in the class should generally be equal to the number of fields reported by the `getFields()` method, unless the `flushContents()` and `initializeContents()` methods are written to combine or split fields. If they are so written, you can define an arbitrary mapping of persistent fields to Java instance fields. For example:

```
class Some {
    int a;
    int b;
    int aPlusb;

    initializeContents(GenericObject, go) {
```



```

        a=go.getField(1, SomeClassInfo);
        b=go.getField(2, SomeClassInfo);
        c=a+b;
    }
    ...
}

```

In a separate file:

```

public class SomeClassInfo
{
    static Field[] fields=
    { field.createInt("a");
      field.createInt("b");
    }
}

```

Persistent-only fields

The list of `Field` objects returned by the `getFields()` method might include one or more fields that are not in the Java class definition. Such fields are persistent-only fields. The `flushContents()` method associated with the class must set the field value in the generic object based on other fields of the class.

Variable initializers

If you manually annotate a class, you should avoid using variable initializers to initialize persistent fields of persistence-capable objects. Instead, perform the initialization in the constructor. This is because the values computed by the variable initializer expression typically are overwritten by the `com.odi.IPersistent.initializeContents()` method. When an object is actually fetched from the database, the fields are initialized with their correct persistent values.

Example

An example of how you might use transient-only and persistent-only fields is in the `demo` directory that is included in `ObjectStore`. In the `rep` example, `Rectangle.a` and `Rectangle.b` are transient-only fields, while `ax`, `ay`, `bx`, and `by` are persistent-only fields. Following is the part of the example that shows this:

```

package com.odi.demo.rep;

/**
 * A Rectangle has two Points, representing its upper-left
 * and lower-right corners. However, its persistent
 * representation is formed by storing the x and y coordinates of
 * the two points, rather than the points themselves. This
 * demonstrates the control that the definer of a persistent
 * class has over the persistent representation. Note that
 * Identity of the Point objects is not preserved, since thePoint
 * objects are not persistent objects. */

import com.odi.*;

public class Rectangle implements IPersistent {

```

```
transient private com.odi.imp.ObjectReference ODIREf;
transient public byte ODIOBJECTState;

transient Point a;
transient Point b;

static ClassInfo classInfo
    = ClassInfo.register(new RectangleClassInfo());

public com.odi.imp.ObjectReference ODIGETRef() {
    return ODIREf;
}

public void ODISETRef(com.odi.imp.ObjectReference objRef) {
    ODIREf = objRef;
}

public byte ODIGETState() {
    return ODIOBJECTState;
}

public void ODISETState(byte state) {
    ODIOBJECTState = state;
}

Rectangle(Point a, Point b) {
    this.a = a;
    this.b = b;
}

void describe() {
    System.out.println("Rectangle with two points:");
    a.describe();
    b.describe();
}

/* Annotations for persistence. */
Rectangle(ClassInfo ignored) {}

public void initializeContents(GenericObject handle) {
    a = new Point(handle.getIntField(1, classInfo),
        handle.getIntField(2, classInfo));
    b = new Point(handle.getIntField(3, classInfo),
        handle.getIntField(4, classInfo));
}

public void flushContents(GenericObject handle) {
    handle.setIntField(1, a.x, classInfo);
    handle.setIntField(2, a.y, classInfo);
    handle.setIntField(3, b.x, classInfo);
    handle.setIntField(4, b.y, classInfo);
}

public void clearContents() {
    a = null;
    b = null;
}
```

```

    }

    /* This class is never used as a persistent hash key. */
    public int hashCode() {
        return super.hashCode();
    }
}

```

In a separate file:

```

public class RectangleClassInfo extends ClassInfo
{
    public IPersistent create() { return new Rectangle(this); }
    public Class getClassDescriptor() throws
        ClassNotFoundException {
        return Class.forName("com.odi.demo.rep.Rectangle");
    }

    public Field[] getFields() { return fields; }
    private static Field[] fields =
    { Field.createInt("ax"),
      Field.createInt("ay"),
      Field.createInt("bx"),
      Field.createInt("by"), };}

```

Defining Persistence-Aware Classes

A persistence-aware class is a class whose instances

- Can operate on persistent objects
- Cannot be stored in a database

For a class to be persistence aware, you must annotate it so that it includes calls to the `ObjectStore.fetch()` and `ObjectStore.dirty()` methods. The `fetch()` method makes the contents of a persistent object available to be read. The `dirty()` method makes the contents of a persistent object available to be modified.

To make a class persistence aware, modify each method that references

- Nonstatic fields of persistence-capable classes
- Array elements of arrays that might be persistent

Modify each method so that it calls the `ObjectStore.fetch()` or `ObjectStore.dirty()` method. This call must be before any attempt to access the contents of the persistent object. The `fetch()` and `dirty()` methods make the contents of persistent objects available.

A persistence-aware class includes the `fetch()` and `dirty()` annotations. It does not include the other annotations that are required for a class to be persistence capable.

Following Postprocessor Conventions

If you plan to define all required annotations explicitly, you need not be concerned with postprocessor conventions. However, if you plan to insert some annotations explicitly and use the postprocessor to insert other annotations, you must follow these postprocessor conventions.

- The name of the `ClassInfo` subclass must have the following format:

`class_nameClassInfo`

For example, if you define the `Boat` class, the name of the associated subclass of `ClassInfo` must be `BoatClassInfo`.

- In the `ClassInfo` subclass definition, when you define the hollow object constructor, it must take a single argument of type `ClassInfo`. See page 246.

Annotating Abstract Classes

Persistence-capable classes and their superclasses, even if they are abstract, must each have a corresponding `ClassInfo` subclass. But an application does not create instances of abstract classes, so you cannot write the required `create()` method in the `ClassInfo` subclass in the usual way. Define the `create()` method so that it returns null. Because this method is never called, it is safe to define it this way.

Now suppose you define the following two classes:

```
abstract class Y {
    int yValue;
    abstract void doSomething();
}

class X extends Y {
    float xValue;
    void doSomething() {}
}
```

Class `Y` must have an associated `ClassInfo` subclass and class `X` must have an associated `ClassInfo` subclass. The `ClassInfo` subclass associated with `X` does not extend the `ClassInfo` subclass associated with `Y`.

In the `ClassInfo` subclass for `X`, the `Field` array must include only those fields defined explicitly in `X`; `XClassInfo.getFields()` must report only the

immediate persistent fields in *X*. The `ClassInfo` subclass for *Y* defines a `Field` array that contains the fields explicitly defined in *Y*.

Creating and Accessing Fields in Annotations

As part of the process of manually defining a class that is persistence capable, the required annotations must, among other things,

- Define an `initializeContents()` method in the persistence-capable class
- Define a `flushContents()` method in the persistence-capable class
- Define a `getFields()` method in the `ClassInfo` subclass

To define these methods correctly, you must know how `ObjectStore` makes persistent objects accessible and the methods that are available to create and access individual fields in an object. To help you do this, this section discusses the following topics:

- Making Persistent Objects Accessible
- Creating Fields
- Getting and Setting Generic Object Field Values
- Methods for Creating Fields and Accessing Them in Generic Objects

Making Persistent Objects Accessible

The `ObjectStore.fetch()` method makes the contents of a persistent object available to be read by an application. The `ObjectStore.dirty()` method makes the contents of a persistent object available to be updated by an application.

To execute a `fetch()` or `dirty()` call, `ObjectStore` first checks whether a `fetch()` or `dirty()` call was already invoked on the object in the current transaction. If it was, `ObjectStore` does nothing and the program continues. If it was not, `ObjectStore` executes the `fetch()` or `dirty()` call, as required.

When `ObjectStore` retrieves a persistent object, it calls the `initializeContents()` method that you defined. The `initializeContents()` method calls methods on `GenericObject` to obtain

the field values for the persistent object. The result is that your program has access to the desired data.

ObjectStore provides the `GenericObject` class for transferring data between a database and a Java application or applet. A generic object represents an object's data as it is stored in the database. A generic object is a temporary buffer that ObjectStore uses while it is copying data from the database into a persistent object or writing data into the database from a persistent object. ObjectStore creates instances of `GenericObject` as needed. You do not define subclasses of `GenericObject`, nor do you create instances of `GenericObject`.

For an object that was not already retrieved, ObjectStore copies the contents of the object from the database into the `GenericObject` instance. It then passes this instance to the `initializeContents()` method defined in the persistence-capable class.

Suppose you called the `dirty()` method on a persistent object and modified it. To update the object in the database, commit the transaction. This causes ObjectStore to create an instance of `GenericObject` to hold the contents of your object. Then ObjectStore calls the `flushContents()` method that you defined when you defined the persistence-capable class.

The `flushContents()` method must call methods on the `GenericObject` instance that store the object's field values in the generic object. ObjectStore calls the `flushContents()` method as needed to copy the new contents of the object into the database.

Creating Fields

ObjectStore provides the `Field` class to represent a Java field in a persistent object. When you define a persistence-capable class, you must define a `getFields()` method in the required `ClassInfo` subclass. This method provides a list of the nonstatic fields (also called instance variables) whose values are being stored and retrieved.

Description of `getFields()`

The `getFields()` method must return an array that contains the nonstatic persistent object fields. The order in which they appear in the array implies their associated field numbers. This array must include only those fields defined in the persistence-capable class and not any inherited fields.

Field numbers represent the position of a nonstatic field within the list of all nonstatic fields defined for the class and its superclasses. The first field has field number 1. (Note that the first field number is not 0.)

Order of fields

When you define the `getFields()` method in the `ClassInfo` subclass, you determine the order and, therefore, the number of each field even though you do not explicitly assign any numbers. `ObjectStore` assigns the numbers according to the order in which the values are returned from the field create methods defined in the `getFields()` method. The field numbers are consecutive with no gaps. For example:

Example

```
public Field[] getFields() {return fields;}

    private static Field[] fields = {
        Field.createString("name"),
        Field.createInt("age"),
        Field.createClassArray("children",
            "com.odi.demo.people.Person", 1)
    };
```

The previous definition causes `ObjectStore` to associate 1 with the `name` field, 2 with the `age` field, and 3 with the `children` field.

When you define the `initializeContents()` and `flushContents()` methods, you must specify the correct field number for each field that the methods get and set.

Creation methods

The `Field` class provides a create method for each Java data type. Minimally, the create methods on the `Field` object

- Return the created `Field` object
- Take a `String` parameter that specifies the name of the field

There are separate create methods for singleton and array fields of each primitive type. There are also string fields, class fields, and interface fields. The complete list of field create methods is in *Methods for Creating Fields and Accessing Them in Generic Objects* on page 256.

Getting and Setting Generic Object Field Values

As described earlier, `ObjectStore` provides the `GenericObject` class to transfer objects between the database and an application. Consequently, when you define a persistence-capable class, you must define the `initializeContents()` method to retrieve values from fields in instances of `GenericObject` and the `flushContents()` method to set values in fields of instances of `GenericObject`.

When you define the `initializeContents()` and `flushContents()` methods, you must use a method that is appropriate for the type of each field in the instance of `GenericObject`. For example, for each character field, you must use the

- `getCharField()` method in the `initializeContents()` method
- `setCharField()` method in the `flushContents()` method

There is a different method for getting and setting each Java type. To get or set an array of any type, you define the `getArrayField()` and `setArrayField()` methods, respectively. In the `initializeContents()` method, be sure to call the methods that get the values. In the `flushContents()` method, be sure to call the methods that set the values. The methods that get and set fields in a generic object are listed in the table in *Methods for Creating Fields and Accessing Them in Generic Objects* on page 256.

Methods for Creating Fields and Accessing Them in Generic Objects

<i>Kind of Java Field</i>	<i>Code</i>	<i>Method That Operates on It</i>
Single byte	<code>byte</code>	<code>Field.createByte()</code> <code>GenericObject.getBytesField()</code> <code>GenericObject.setBytesField()</code>
Array of bytes	<code>byte[]</code>	<code>Field.createByteArray()</code> <code>GenericObject.getArrayField()</code> <code>GenericObject.setArrayField()</code>
Single character	<code>char</code>	<code>Field.createChar()</code> <code>GenericObject.getCharField()</code> <code>GenericObject.setCharField()</code>
Array of characters	<code>char[]</code>	<code>Field.createCharArray()</code> <code>GenericObject.getArrayField()</code> <code>GenericObject.setArrayField()</code>
Single 16-bit integer	<code>short</code>	<code>Field.createShort()</code> <code>GenericObject.getShortField()</code> <code>GenericObject.setShortField()</code>
Array of 16-bit integers	<code>short[]</code>	<code>Field.createShortArray()</code> <code>GenericObject.getArrayField()</code> <code>GenericObject.setArrayField()</code>
Single 32-bit integer	<code>int</code>	<code>Field.createInt()</code> <code>GenericObject.getIntField()</code> <code>GenericObject.setIntField()</code>
Array of 32-bit integers	<code>int[]</code>	<code>Field.createIntArray()</code> <code>GenericObject.getArrayField()</code> <code>GenericObject.setArrayField()</code>

<i>Kind of Java Field</i>	<i>Code</i>	<i>Method That Operates on It</i>
Single 64-bit integer	long	Field.createLong() GenericObject.getLongField() GenericObject.setLongField()
Array of 64-bit integers	long[]	Field.createLongArray() GenericObject.getArrayField() GenericObject.setArrayField()
Single 32-bit floating-point number	float	Field.createFloat() GenericObject.getFloatField() GenericObject.setFloatField()
Array of 32-bit floating-point numbers	float[]	Field.createFloatArray() GenericObject.getArrayField() GenericObject.setArrayField()
Single 64-bit floating-point number	double	Field.createDouble() GenericObject.getDoubleField() GenericObject.setDoubleField()
Array of 64-bit floating-point numbers	double[]	Field.createDoubleArray() GenericObject.getArrayField() GenericObject.setArrayField()
Single Boolean value	boolean	Field.createBoolean() GenericObject.getBooleanField() GenericObject.setBooleanField()
Array of Boolean values	boolean[]	Field.createBooleanArray() GenericObject.getArrayField() GenericObject.setArrayField()
Single string value	String	Field.createString() GenericObject.getStringField() GenericObject.setStringField()
Array of string values	String[]	Field.createStringArray() GenericObject.getArrayField() GenericObject.setArrayField()
A class		Field.createClass() GenericObject.getClassField() GenericObject.setClassField()
A class array		Field.createClassArray() GenericObject.getArrayField() GenericObject.setArrayField()

<i>Kind of Java Field</i>	<i>Code</i>	<i>Method That Operates on It</i>
An interface		Field.createInterface() GenericObject.getInterfaceField() GenericObject.setInterfaceField()
An interface array		Field.createInterfaceArray() GenericObject.getArrayField() GenericObject.setArrayField()

Chapter 10

Controlling Concurrency

This chapter provides information about ways that you can control concurrency. The APIs described in this chapter make it easier for your application to access data and less likely that your application must wait to access that data. In addition, you can choose to limit access by other users to the same data.

Contents	This chapter discusses the following topics:	
	Reducing Wait Time for Locks	259
	Using Multiversion Concurrency Control (MVCC)	261
	Checkpoint: Committing and Continuing a Transaction	264
	Locking Objects, Clusters, Segments, and Databases to Ensure Access	267
	Installing Schema Information in Batch Mode	271

Reducing Wait Time for Locks

What can you do to reduce the overhead of waiting for locks? One application can reduce the waiting overhead for other concurrent applications by avoiding locking data unnecessarily and by avoiding locking data for unnecessarily long periods of time. This section describes several techniques for minimizing wait time.

Clustering

One way to help avoid locking data unnecessarily involves not clustering objects together when they are not normally accessed together. Suppose that during a given transaction, an application requires `object-a` but not

`object-b`. The two objects are stored on the same page. Because ObjectStore performs page-level locking, when you access one of these objects, ObjectStore locks both of them. This prevents other processes from accessing either object until the end of the transaction. If you store `object-b` in a different cluster from `object-a`, you guarantee that the objects are on different pages. Therefore, the objects are not locked together.

Transaction Length

Making transactions shorter is one way to avoid locking data for unnecessarily long periods of time. If you do this, you must still ensure that objects in the database are in a consistent state between transactions.

The disadvantage of using shorter transactions is that it can mean using a greater number of transactions. This can increase network overhead because each transaction commit requires the client to send a *commit message* to the server. Nevertheless, this extra network overhead is often outweighed by the savings from shorter waits for locks to be released.

It is sometimes particularly important to make transactions that store new objects in the database or destroy persistent objects as short as possible. This is because many write locks are required, which can decrease the level of concurrency.

Multiversion Concurrency Control (MVCC)

Read-only transactions can use *multiversion concurrency control*, or MVCC. With MVCC, an application can perform nonblocking reads of a database. This allows another application to update the database concurrently, with no waiting by either the reader or the writer. See Using Multiversion Concurrency Control (MVCC) on page 261.

Lock Timeouts

Lock timeouts provide the ability to limit the time that ObjectStore waits to obtain a lock. When you try to obtain a lock on an object, segment, or database, you can specify the number of milliseconds for which it is all right to wait for the lock. See Locking Objects, Clusters, Segments, and Databases to Ensure Access on page 267.

Conflicts Caused by Schema Installation

If you find that there are concurrency conflicts caused by incremental schema installation, you can install schema information in batch mode. See *Installing Schema Information in Batch Mode* on page 271.

Using Multiversion Concurrency Control (MVCC)

When an application uses MVCC, it can perform nonblocking reads of a database. This means that another ObjectStore application can concurrently update the database. Neither the reader nor the writer has to wait for the other. To use MVCC, specify `ObjectStore.MVCC` or `ObjectStore.MULTI_DB_MVCC` as the open type when you open a database.

When Is MVCC Appropriate?

MVCC is useful when your application contains a transaction that

- Does not modify a database
- Does not require a view of the database that is completely up to date, but can instead rely on a snapshot of the data
- Does not depend on having data in a database opened for MVCC be transaction consistent with data in other databases

How Does MVCC Work?

In each transaction in which an application accesses a database opened for MVCC, it is as if the application were viewing a snapshot of the database. This snapshot

- Is taken sometime during the transaction
- Is internally consistent
- Might not contain changes that were committed by other sessions during the transaction
- Contains all changes that were committed before the transaction started

If an application accesses databases that have been opened for multidatabase MVCC it will see snapshots of those databases taken at the same moment in time.

Obtaining Read Locks

When an application has a database opened for MVCC, the application never has to wait for read locks on the database. When an application reads data from a database opened for MVCC, the application never causes other applications to wait for write locks. In addition, when an application accesses a database it has opened for MVCC, the application never causes a deadlock.

Accessing Multiple Databases in a Transaction

You can open databases for MVCC in either single-database or multidatabase mode. When an application reads a database opened for single-database MVCC, the snapshot it views is internally consistent but potentially out of date. This means that the snapshot might not be consistent with other databases accessed in the same transaction. Even two databases, both of which are opened for MVCC, might not be consistent with each other. Updates might be performed on one of the databases in between the times of their snapshots.

When an application reads databases opened for multidatabase MVCC, the snapshot is a consistent, aggregate view of all data items in all the databases at the time the snapshot was taken.

Serializability

A snapshot might be out of date by the time an application reads some data. However, if each transaction that accesses a database opened for MVCC accesses only that one database, MVCC retains serializability. Such a transaction views a database state that would have resulted from some serial execution of all transactions. All transactions produce the same effects as would have been produced by the serial execution.

Opening a Database for MVCC Access

To use MVCC, specify `MVCC` or `MULTI_DB_MVCC` as the open mode when you open a database. For example, to open a database in single-database MVCC:

```
Database db = Database.open(
    "myDb.odb", ObjectStore.MVCC);
```

After an application opens a database for MVCC, it can read that database without waiting for locks or blocking other applications.

You cannot change the open mode of an open database. To change the type of access to a database, you must close the database, then reopen it with a specification of the new open mode.

If you try to update data in a database that you opened for MVCC, `ObjectStore` signals `UpdateReadOnlyException`.

In a session, all cooperating threads use the same access mode for a particular database. For example, a thread cannot open a database for update if a cooperating thread has already opened that database for MVCC. However, if a thread opens a database for MVCC, a thread in another session, that is, a noncooperating thread, can open the same database for update.

In the same transaction, your application can open one or more databases for MVCC and open other databases for read-only or update.

Determining Whether a Database Is Opened for MVCC

To determine whether a database is opened for MVCC, call the `Database.getOpenMode()` method. The method signature is

```
public int Database.getOpenMode()
```

If the database is opened for single-database or multidatabase MVCC, `ObjectStore` returns the `ObjectStore.MVCC` or `ObjectStore.MULTI_DB_MVCC` constant, respectively. Otherwise, if the database is open, `ObjectStore` returns either the `ObjectStore.UPDATE` or `ObjectStore.READONLY` constant.

When `ObjectStore` opens a database as a result of following a cross-database pointer, the automatic open mode can be `ObjectStore.MVCC` or `ObjectStore.MULTI_DB_MVCC`.

Updating the Snapshot

In a transaction in which you access a database that you opened for MVCC, you might want to update the snapshot periodically. If several databases are open for MVCC, updating the snapshot updates the data in all the databases open in that mode. Updating the snapshot can be done in the following two ways:

- End the transaction with `commit()` or `abort()` and start a new transaction.
- Call `checkpoint()` on the transaction. This has the effect of committing the transaction and starting a new transaction, but it does not incur the

overhead of a new transaction. See *Checkpoint: Committing and Continuing a Transaction* on page 264.

To help you decide when to do this, you can find out whether, during your transaction, there were any write locks on the database you opened for MVCC. A write lock indicates that another application might have made a modification to the database. To do this, call the `Transaction.hasLockContention()` method on your in-progress transaction. The method signature is

```
public boolean hasLockContention()
```

`ObjectStore` returns `true` if a server involved in your transaction might have write locked an object that was read by your application. PSE Pro always returns `false`.

With a return value of `true`, if you commit your transaction and start a new one, or checkpoint your transaction, you might have access to updated data. However, you might also have access to the same data if the application that had the write lock either aborted the transaction or did not commit any changes.

Where to Find Additional Information

Additional information about MVCC can be found in the `ObjectStore C++` documentation. See the *Advanced C++ API User Guide* for information about

- MVCC and the transaction log
- Conflict detection
- Propagating data from the log to the database
- Transaction locking examples

Checkpoint: Committing and Continuing a Transaction

With the `Transaction.checkpoint()` method, you get the effect of committing a transaction, then continuing work in a new transaction in which you have read locks on all or most of the persistent objects that were locked in the committed transaction. This is useful when

- You are making modifications to a database. You want to periodically commit your changes but continue updating the database without

intervention. For example, you might be loading new data into the database.

- You want to make your changes available to MVCC readers.
- You opened a database for MVCC and you want an updated snapshot.

Note This checkpoint differs from a conventional checkpoint. In this checkpoint, an application might not have all the locks after the checkpoint that it had before the checkpoint. The details are explained in the next section.

When to checkpoint Before you checkpoint a transaction, you must ensure that the database is in a consistent state because the state is made persistent.

Caution You should be aware of the following issues when you are using the `Transaction.checkpoint()` method:

- If your application checkpoints a transaction while an annotated method is executing, your program might incorrectly access persistent objects after the checkpoint. For more information about this and a work around, see [Troubleshooting Access to Persistent Objects](#) on page 148.
- During the checkpoint, you must ensure that no other thread tries to access the database.

Method Signatures The `checkpoint()` method has two overloadings. The first overloading takes no argument. The method signature is

```
public void checkpoint()
```

The second overloading has an argument that specifies the state of persistent objects after a checkpoint operation. The method signature is

```
public void checkpoint(int retain)
```

The value of `retain` can be one of the following:

- `ObjectStore.RETAIN_STALE` resets the contents of persistent objects to default values and makes all persistent objects stale.
- `ObjectStore.RETAIN_HOLLOW` resets the contents of persistent objects to default values but makes the persistent objects hollow. Before and after the checkpoint, you can use references to the same objects.

Advantages of a Checkpoint

The advantage of a checkpoint is that there is less overhead than when you actually end one transaction and start another. When you checkpoint a transaction, it is as if you committed the transaction, then immediately

started a new transaction. However, in the new transaction, you already have read locks on most or all of your persistent objects.

If another session is waiting for a write lock on a persistent object that was locked in your transaction, you lose that lock when you checkpoint the transaction. As long as another session is not waiting for a write lock on an object that was associated with your transaction, you reacquire as read locks any locks you had before the checkpoint.

After the checkpoint, the persistent objects are stale or hollow, according to the one you specify when you call `checkpoint()`. If you specify that objects should be hollow, you do not have to start from a root object to set up your access to objects. Your application's access to objects is the same before and after the checkpoint.

After a checkpoint, `ObjectStore` has read locks on the same objects as before the checkpoint, unless another session was waiting for a write lock on one of these objects. In that case, your transaction loses the lock.

If there were any write locks before the checkpoint, `ObjectStore` changes them to read locks or gives them to any sessions waiting for those write locks. Consequently, you might have to wait for locks or you might get a deadlock when you try to update the database again.

Suppose your application calls `Transaction.checkpoint()`, then the transaction started by the `checkpoint()` method is aborted. `ObjectStore` does not commit any changes to the database that were made after the checkpoint operation. Any changes made before the checkpoint remain committed.

Setting a Default Checkpoint Retain State for a Session

You can use the following two methods to set the default retain state for persistent objects after a checkpoint in a transaction:

- `Transaction.setDefaultCommitRetain()`
- `Transaction.setDefaultRetain()`

The default retain state set by these methods is in effect for the duration of the session in which they are called.

Use `setDefaultCommitRetain()` to specify a default retain state for persistent objects when `Transaction.commit()` is called without a *retain* argument. The `setDefaultCommitRetain()` method also sets the default

retain state for persistent objects when `Transaction.checkpoint()` is called without a *retain* argument. The method signature is

```
public void setDefaultCommitRetain(int retain)
```

Use `setDefaultRetain()` to specify a default commit retain state for persistent objects when `Transaction.commit()` is called without a *retain* argument. This method also sets the default retain states for persistent objects, when `Transaction.abort()` and `Transaction.checkpoint()` are called without a *retain* argument. The method signature is

```
public void setDefaultRetain(int retain)
```

Note

If you specify a *retain* value for `setDefaultRetain()` or `setDefaultCommitRetain()` that is *not* invalid for `Transaction.checkpoint()` (such as `RETAIN_UPDATE`, `RETAIN_TRANSIENT`, and in some cases `RETAIN_READONLY`), then `ObjectStore` automatically maps these *retain* values to `RETAIN_HOLLOW`.

Also, when you are using either method to set a default retain state, the method that was last called overrides any default retain state that was set previously. For example, if an application calls `setDefaultCommitRetain()` first and then calls `setDefaultRetain()`, the retain state for `checkpoint()` and `commit()` is specified by the second call.

Setting Persistent Objects to a Default State

To checkpoint a transaction and to set the state of persistent objects to the state specified by `Transaction.setDefaultCommitRetain()` or by `Transaction.setDefaultRetain()`, call the `checkpoint()` method without any arguments. For example:

```
tr.checkpoint();
```

Locking Objects, Clusters, Segments, and Databases to Ensure Access

There are times when you want to ensure your own access to objects and limit access to those objects by other sessions. This can be when you want to ensure that

- An operation completes without interruption.
- All objects are immediately available.

- There are no deadlocks.

To ensure your own access, you can lock an object, a cluster, a segment, or a database. This means that

- A transaction from another session cannot block your transaction from updating the locked object.
- If you update the object, no other sessions can read the object unless they use MVCC. The use of MVCC by other sessions prevents them from being blocked and from running into a deadlock.

Advantages of locking

The overhead for locking objects is far less than the overhead for reading a database with MVCC. Also, if you want to, you can update a locked object. However, if you lock many individual objects, the overhead might be comparable.

Disadvantages of MVCC

The use of MVCC has some disadvantages:

- You cannot update the database.
- The overhead for MVCC grows with each concurrent update.
- It is possible for multiple databases to be inconsistent with each other.

Note

The `com.odi.useDatabaseLocking` property is a PSE Pro feature. If you are using PSE Pro as well as OSJI, beware of confusing this property with the `OSJI Database.acquireLock()` method. This method allows a session to explicitly lock a particular database for exclusive use.

PSE Pro

The `acquireLock()` methods do nothing in PSE Pro.

Description of Acquire Lock Methods

The following methods allow you to acquire locks on these objects:

- `ObjectStore.acquireLock()` obtains a lock on a specified object. The method signature is

```
public static void ObjectStore.acquireLock(
    Object object, int lockType, int timeoutMillis);
```

- `Cluster.acquireLock()` obtains a lock on a specified cluster. This locks all the objects in the cluster. The method signature is

```
public static void Cluster.acquireLock(
    int lockType, int timeoutMillis);
```

- `Segment.acquireLock()` obtains a lock on a specified segment. This locks all the objects in the segment. The method signature is

```
public static void Segment.acquireLock(
```

```
int lockType, int timeoutMillis);
```

- `Database.acquireLock()` obtains a lock on a database. This locks all the segments, which locks all the objects in the segments. The method signature is

```
public static void Database.acquireLock(
    int lockType, int timeoutMillis);
```

Note

A deadlock can occur if two different sessions call `ObjectStore.acquireLock()` to get a lock at the same time.

Locking Objects for Read or Write Access

The `lockType` parameter indicates whether you want to read or update the locked objects. You must specify one of the following:

- `ObjectStore.READONLY` instructs `ObjectStore` to make the contents of the locked objects available to be read. While you have this read lock, other sessions can obtain read locks as usual.
- `ObjectStore.UPDATE` instructs `ObjectStore` to make the contents of the locked objects available to be modified. You must be in an update transaction and the database must be opened for update. While you have this write lock, no other sessions can access the object unless they use MVCC to do so.

In a transaction, you can reissue the `acquireLock()` call to change the `lockType`.

Specifying the Wait Time for a Lock

If the lock is not available, the `timeoutMillis` parameter indicates the number of milliseconds you are willing to wait for the lock. You can specify

- `ObjectStore.WAIT_FOREVER` to wait until the lock is available.
- 0 if you do not want to wait at all.
- A positive number to indicate the number of milliseconds it is all right to wait. `ObjectStore` rounds up to the nearest number of seconds.

If `ObjectStore` cannot acquire the lock, either because you do not want to wait or because the waiting period has been exceeded, `ObjectStore` signals `LockTimeoutException`.

Releasing Locks

To release locks, you must end the transaction in which you acquired them. `Transaction.checkpoint()` also releases locks, however you reacquire locks as read locks if no other session is waiting for a write lock for the objects you had locked.

Locking Peer Objects

When you lock a Java peer object that identifies a persistent C++ object, `ObjectStore` locks the entire object. However, `ObjectStore` does not lock subobjects that are logically part of the peer object. For example, when you lock a `com.odi.coll` collection, you do not lock the contents of the collection.

You cannot lock Java peer objects that represent transient C++ objects.

Obtaining Information About Concurrency Conflicts

Instances of the `LockTimeoutBlocker` class represent clients that have been involved in concurrency conflicts.

Client list	To obtain a list of such clients, call the <code>LockTimeoutException.getBlockers()</code> method. The signature is <pre>public LockTimeoutBlocker[] getBlockers()</pre>
Lock type	To find out whether a concurrency conflict was for a read lock or a write lock, call <code>LockTimeoutException.getLockType()</code> . This returns <code>ObjectStore.READONLY</code> or <code>ObjectStore.UPDATE</code> . The method signature is <pre>public int getLockType()</pre>
Client names	To obtain the name of the client that caused the conflict, call <code>LockTimeoutBlocker.getApplicationName()</code> . The method signature is <pre>public String getApplicationName();</pre>
Process IDs	To obtain the process ID of the process running a client that caused a conflict, call the <code>LockTimeoutBlocker.getPID()</code> method. The method signature is <pre>public int getPID()</pre>
Host names	To obtain the name of the host for the process that was running the client that caused the conflict, call the <code>LockTimeoutBlocker.getHostName()</code> method. The method signature is <pre>public String getHostName()</pre>

Setting the Client Name

To allow other applications to see meaningful client names when ObjectStore signals `LockTimeoutException`, your application should set a name for itself. Use the system property `com.odi.applicationName` to do this. Specify the name of the application for the current client. Include it in the list of properties that you specify when you create a session.

Note

The application name is set by the first session that specifies the `com.odi.applicationName` property. This property is ignored if you specify it when you create a later session.

Helping Determine the Transaction Victim in a Deadlock

When there is a deadlock, the server must choose a transaction to abort. The server uses several criteria to pick the victim. One of the criteria is the transaction priority, which you can set with the `Transaction.setPriority()` method.

To obtain the priority that is assigned to transactions started by the current session, call the `Transaction.getPriority()` method.

Every client has a transaction priority, which is a value in the range 0 to 0xffff. The default value is 0x8000, which is in the middle of the range. The value 0 has a special meaning.

In a deadlock situation, the ObjectStore server compares the transaction priorities of clients involved in the deadlock. If the lowest transaction priority is held by only one client, this client is the victim. If the lowest transaction priority is held by more than one client, the server chooses a victim according to the setting of the `Deadlock Victim` server parameter. You can find information about this parameter in *Managing ObjectStore*.

If all transactions in a deadlock have a transaction priority of 0, the server aborts all of them. While this is not a useful way to run a program, it is useful for debugging. You can run several clients under debuggers and have them all set their priorities to 0. When a deadlock happens, all of them abort and you can see what each of them was doing. You should use a priority of 0 only if you want this special debugging behavior.

Installing Schema Information in Batch

Mode

Sometimes, schema installation causes concurrency conflicts. If it does, you can minimize this by installing schema in batch mode.

By default, ObjectStore stores schema information in databases incrementally as needed. If you want to, you can instruct ObjectStore to install schema information in batch mode before it is actually needed. This section provides information about how to do that. The topics discussed are

- Background About Schema Information
- Procedure for Installing Schema in Batch
- Identifying the Application Types
- Creating a Database with Batch Schema Installation
- Installing Application Types in the Database Schema
- If You Do Not Run the Postprocessor

Background About Schema Information

Schema information describes the classes of objects that are stored in the database. ObjectStore stores schema information in each database.

ObjectStore supplies schema information for the internal C++ types that it requires. If your application accesses C++ classes, the process of building your Java/C++ application makes additional C++ schema information available. ObjectStore also needs schema information for classes you define and for any third-party libraries that your classes use. The postprocessor or ObjectStore run time creates this schema information as Java `ClassInfo` subclasses that are translated into C++ schema information as needed.

ObjectStore dynamically stores schema information in each database as needed. This is referred to as incremental schema installation. If you want to, you can instruct ObjectStore to install internal schema information when it creates the database and application schema information when you invoke a method to do so. In this way, you install schema information before it is needed. This is referred to as batch schema installation.

The advantages of batch schema installation are

- You minimize the chance of a concurrency conflict caused by schema installation.

- It is more efficient when a large percentage of the schema information is actually needed.

The disadvantages are

- It takes a little longer to create the database.
- An initial database is larger than an initial database that uses incremental schema installation.
- Schema information that your application never uses might be installed in a database.

Procedure for Installing Schema in Batch

To perform batch schema installation for a database, follow these steps:

- 1 Use the postprocessor to create one or more objects that identify the application types.
- 2 Create a database with the overloading of `Database.create()` that allows you to specify batch schema installation.
- 3 Start an update transaction.
- 4 Install the application types in the database. These are the types identified in the object or objects created in step 1.

In the unusual situation in which you do not run the postprocessor, you follow a different procedure to perform batch schema installation. See *If You Do Not Run the Postprocessor* on page 277.

Identifying the Application Types

The application types are the persistence-capable classes you define and any persistence-capable classes in third-party libraries that your classes refer to. To install schema information in batch mode, you need to create one or more objects that identify all application types. There are two postprocessor options that you can specify to do this.

Individual types

When you run the postprocessor to make classes you define persistence capable, specify the `-summary` option. The format is

```
-summary gen_class_name
```

This option instructs the postprocessor to generate a class with the name `gen_class_name`. This generated class extends the `PersistentTypeSummary` class. The `gen_class_name.getPersistentClasses()` method returns a

list of the classes that were made persistence capable in this execution of the postprocessor.

The generated class has a no-arguments constructor that passes arrays to the `PersistentTypeSummary` class constructor.

Types in libraries

Classes in libraries that have already been postprocessed must have an associated generated class (an associated summary) that identifies the persistence-capable classes in the library. You must do one of the following:

- Create this summary yourself when you specify the `-summary` option during postprocessing of the library.
- Receive this summary with its associated library from a third-party vendor.
- Manually code the summary yourself by defining a class that extends the `PersistentTypeSummary` class.

You can create one object that contains all relevant summaries, or you can have two or more summaries, which collectively identify all application types. Technical Support recommends that you include the summaries for libraries in the summary the postprocessor creates for the classes that it makes persistence capable. To do this, run the postprocessor and specify the `-includesummary` option for each library summary. The format for specifying this option is

```
-includesummary inc_class_name
```

Replace `inc_class_name` with the name of a class generated by a previous execution of the postprocessor with the `-summary` option. You must know the name of the generated class. If you did not postprocess the library yourself, you must get the name of the generated class from the library vendor.

When you specify the `-includesummary` option, you must also specify the `-summary` option. The postprocessor includes the specified summaries in the new summary it creates. It does not matter whether or not the postprocessor is also annotating any classes.

Library providers

If you are providing a library that contains persistence-capable classes, you must also provide the summary object that identifies those classes. `ObjectStore` uses the summary object to install schema information for your library at run time. This allows users to install a new version of a library without having to rebuild the application type summary.

General use

For a given execution of the postprocessor,

- You can specify the `-summary` option either once or not at all.
- You can specify the `-includesummary` option zero, one, or more times.
- If you specify the `-includesummary` option, you must also specify the `-summary` option.

Creating a Database with Batch Schema Installation

When you create a database, you determine whether you can perform batch schema installation. To allow batch installation, use the following overloading of `Database.create()`:

```
public static Database create(
    String name, int fileMode, int schemaInstallMode);
```

The `name` parameter specifies the pathname of a file. The path can specify a relative name, fully qualified name, remotely mounted directory, or host name and host-relative pathname. An ObjectStore server must be available for the directory that contains the specified file.

The `fileMode` parameter specifies the access mode for the database.

The `schemaInstallMode` parameter specifies the schema installation mode for the database. The value of this parameter must be either of the following:

- `ObjectStore.INSTALL_SCHEMA_BATCH` instructs ObjectStore to store schema information immediately for all types other than your application types in the database being created.
- `ObjectStore.INSTALL_SCHEMA_INCREMENTAL` indicates that ObjectStore should install schema incrementally. ObjectStore installs some schema information immediately. When you store an object in the database, if the required schema information is not already in the database, ObjectStore stores it automatically at that time.

Incremental schema installation is the default. When you use the overloading of `Database.create()` that does not include a specification for `schemaInstallMode`, it is as if you specified `ObjectStore.INSTALL_SCHEMA_INCREMENTAL`.

After you create a database, you cannot change the schema installation mode.

Note

If you use Java/C++ interoperability and you specify batch schema installation when you create a database, you must not use C++ to change the schema installation flag. Doing so might cause concurrency conflicts during future access to the database.

Installing Application Types in the Database Schema

If you have one or more `PersistentTypeSummary` objects that identify your application types, you can install schema information using batch mode in that database. To do so, ensure that the database is opened for update, start an update transaction, then invoke the `Database.installTypes()` method. The signature is

```
public void installTypes(PersistentTypeSummary summary);
```

Summary parameter

The `summary` parameter must be an object that identifies the persistence-capable types used by the application. If there are multiple objects that collectively identify all application types, you must do either of the following:

- Invoke `installTypes()` for each summary object.
- Construct a summary object that identifies these individual summary objects as included libraries. Use this new summary object as the parameter to the `installTypes()` method.

Typically, you create the summary object with the `-summary` and `-includesummary` options to the postprocessor. In unusual circumstances, you might explicitly create a summary object with the `com.odi.PersistentTypeSummary` constructor. Information about doing this is in the next section.

Example

For example, suppose you ran the postprocessor with the following command:

```
osjcfp -dest .\osjcfpout Class1 Class2 -summary comcom.xyz.MySummary
```

In your program, create a database with batch schema installation and invoke `installTypes()` with an instance of `MySummary`:

```
Database db = Database.create(
    "mydb.odb",
    ObjectStore.ALL_WRITE | ObjectStore.ALL_READ,
    ObjectStore.INSTALL_SCHEMA_BATCH);
Transaction.begin(ObjectStore.UPDATE);
db.installTypes(new com.xyz.MySummary());
Transaction.current().commit();
```

ClassInfo must be registered

The `installTypes()` method ensures that an instance of the `ClassInfo` subclass associated with each identified persistence-capable class is registered properly. This normally is done automatically by the postprocessor. If `ObjectStore` cannot find the registered `ClassInfo` subclass instance for a class, `ObjectStore` signals `ClassNotRegisteredException`.

When `ObjectStore` signals this exception, it does not install any schema information for any classes. Thus, even if it can find the schema information for 99 out of 100 classes, it does not install any schema information at all if it cannot find schema information for one class.

If `ObjectStore` detects a problem in the type summary, it signals `InvalidSummaryException`.

Indirectly identified classes

`ObjectStore` installs schema information for persistence-capable classes that are directly and indirectly identified. That is, if a class contains a reference to a class that is not identified in the summary, `ObjectStore` tries to install the schema information for the referenced class. If the referenced class refers to another class that is not in the summary, `ObjectStore` tries to install the schema information for that class.

`ObjectStore` installs schema information for any superclasses of specified classes.

Omitting batch installation

Suppose you specify batch schema installation when you create a database, but you forget to install the application types before you store objects in the database. `ObjectStore` behaves as though the incremental schema installation flag were set. When you store the objects, it loads any needed schema information.

You can invoke `installTypes()` at any time and `ObjectStore` installs schema information for application types if that information is not already in the database.

Omitting a type from the summary

Now suppose you forget to identify a persistence-capable class in the summary that you pass to `installTypes()`. When you store an object of this class in the database, `ObjectStore` dynamically stores the schema information for that class.

If You Do Not Run the Postprocessor

In the unusual circumstance that you manually annotate your code instead of running the postprocessor, there is a manual way to create a summary of your application's persistence-capable classes.

The `com.odi.PersistentTypeSummary` class allows your application types to be identified. To create your own summary, invoke the `PersistentTypeSummary` constructor. The signature is

```
public PersistentTypeSummary(
    String[] persistentClasses,
    String[] includedLibrarySummaries);
```

The `persistentClasses` parameter must be an array of names of classes that are persistence capable. This parameter can be null. The `PersistentTypeSummary.getPersistentClasses()` method returns an array of class names that have been identified as persistence capable.

The `includedLibrarySummaries` parameter must be an array of names of classes that extend the `PersistentTypeSummary` class. Typically, these classes identify the persistence-capable classes in a library. This parameter can be null. The

`PersistentTypeSummary.getIncludedLibrarySummaries()` method returns an array of the class names that contain summaries of persistence-capable classes in libraries.

When you define a class that extends the `PersistentTypeSummary` class, it must have a no-argument constructor.

Suppose you identify a persistence-capable class in a summary, but that class is not, in fact, persistence capable. `ObjectStore` signals `ClassNotRegisteredException` at run time if it recognizes this when you invoke the `installTypes()` method on the database.

Chapter 11

Using the Notification Facility

The notification facility allows an application to notify one or more sessions that an event has taken place. Each notification is associated with a location in a database. Your application determines what constitutes an event. In general, an event is anything you want your application to notify other sessions about. For example, a modification to a particular object in a database can be an event.

Contents	This chapter discusses the following topics:	
	Background About How Notification Works	279
	Creating Notifications	282
	Subscribing to Receive Notifications	284
	Sending Notifications	285
	Retrieving Notifications	286
	Reading Notifications	287
	Managing the Notification Process	287

Background About How Notification Works

This section provides information about how the notification system works. It covers these topics:

- What Is a Notification?

- What Is the Flow of a Notification?
- Threads and Notifications
- Transactions and Notifications
- Security

What Is a Notification?

A notification is an ordinary transient Java object. A notification always specifies

- A location. The location can be a persistent object, a cluster, a segment, or a database.
- An integer value. This is the value of the `kind` argument in the `Notification` constructor.
- Information about the event. This is either a message string or an array of bytes. In the `Notification` constructor that takes only two arguments, `ObjectStore` uses a value of null for the third argument, which specifies the message string or byte array.

When a session subscribes to receive notifications for a segment or database, the subscribing session receives any notifications for objects in the cluster, segment, or database.

What Is the Flow of a Notification?

An application creates a notification. Sessions that want information about events related to a particular location subscribe to notifications that specify the location of interest.

When there is an event that involves the location in a notification, the application uses the notification API to send the notification to the `ObjectStore` server.

When the server receives a notification, it determines the sessions that are subscribed for that notification. The server then queues messages to be sent to the receiving sessions. The server returns the number of messages queued and then asynchronously sends notifications to the cache manager of the receiving sessions (the subscribers).

There is a queue inside the cache manager for each session on that host. When the cache manager receives a notification from the server, it puts the notification into the queue of each of the subscribing sessions.

A session that subscribes to one or more notifications usually dedicates a thread to receive notifications. When this thread finds a notification in the cache manager's queue for that session, it removes the notification from the queue and returns it to the associated session, which performs an application-specific action.

Threads and Notifications

A session starts a thread whose sole purpose is to receive notifications. This thread calls `Notification.receive()` with an argument that specifies the length of time to wait for notifications. Each session has its own dedicated thread.

When a session receives a notification, it performs an application-specific action. For example, it might post a Windows message, modify the application's transient data structures, or otherwise queue the notification for processing by another thread. It then waits for the next notification.

The thread dedicated to receiving notifications typically does very little work. It might do queue management, for example, maintaining a priority queue of notifications for another thread, or coalescing similar notifications. However, processing should be minimal so the cache manager notification queue does not overflow. Queue overflow can happen if many notifications arrive in quick succession. The thread receiving the notifications might not be able to keep up with the process of removing the notification from the queue and returning it to the session.

In contrast to most other ObjectStore APIs, `Notification.receive()` is not locked out when other threads are in ObjectStore operations. If the thread does not access persistent data or call other ObjectStore APIs, it can run entirely asynchronously.

Transactions and Notifications

An application can send a notification

- Immediately
- When the transaction commits

Transactions are independent of immediate notifications, subscriptions, unsubscriptions, and notification retrieval.

Sending commit-time notifications, however, is closely integrated with transactions. ObjectStore queues commit-time notifications inside a transaction and sends them when the transaction commits. If the transaction

aborts, the application never sends the commit-time notifications. This is useful when you want dispatch of the notification to be contingent on a database modification that becomes visible when a transaction commits.

There are no restrictions on transaction types. The enclosing transaction can be read-only or update. Databases can be opened read-only, update, or MVCC.

As always, database changes made by a session are not visible to other sessions until the transaction commits. Therefore, all notifications that indicate changes to persistent data should be made at commit time.

Security

To send or subscribe to notifications, a session must open the database that contains the referenced location. If a session does not open a database, it cannot send or receive notifications associated with that database. If you do not have permission to open a database, you cannot send or subscribe to notifications on objects in that database.

Within a database, notifications are not integrated with ObjectStore security. A session can subscribe to notifications and send notifications that reference database locations in any segment.

Creating Notifications

When an application creates a notification, the database that contains the referenced object must be open. It does not matter whether a transaction is in progress. However, an object that an application passes to a `Notification` method cannot be a stale object.

Descriptions of Constructors

The constructors for creating notifications have the following signatures:

- `public Notification(Object location, int kind)`
- `public Notification(Object location, int kind, byte[] data)`
- `public Notification(Object location, int kind, String message)`

The `location` parameter specifies a persistent object. It indicates the location at the beginning of the object. You must specify a persistent object. It is not

sufficient for the *location* object to be persistence capable. If the specified object is not persistent when you try to construct a notification, `ObjectStore` signals `ObjectNotPersistentException`. To avoid this, call the `ObjectStore.migrate()` method to store the object in the database before you create the notification.

You can specify a Java peer object if it identifies a persistent C++ object. See Chapter 3, Writing the Application in *Developing Java Applications That Access C++*, .

The *kind*, *data*, and *message* parameters provide information about the event. Every notification has a *kind* parameter. If you specify a negative argument, `ObjectStore` signals `IllegalArgumentException`.

If you do not want to attach a message or data to the notification, specify null for the third argument when you create the notification, or do not specify a third argument. For example, the following two code fragments do the same thing:

```
String a = null;
new Notification(foobar, 102, a);

new Notification(foobar, 102);
```

If you do specify a third argument, it is a sequence of bytes. For convenience, `ObjectStore` allows you to pass in a Java `String` instead of a sequence of bytes. `ObjectStore` uses UTF-8 encoding to encode *message* arguments into a sequence of bytes.

When `ObjectStore` sends a notification, it makes the *kind* parameter and the *data* or *message* parameter, if there is one, available to subscribers. If a notification includes a null string ("null"), it is received as an empty string ("").

Retaining References to Persistent Objects

A notification always contains a reference to a persistent object. You can create a notification in one transaction, then use it in a subsequent transaction or between transactions.

To do so, you must ensure that the reference is not to a stale object. If it is, `ObjectStore` signals `ObjectException`.

To ensure a nonstale object, you can evict the referenced object or commit or abort the transaction with a retain type other than `ObjectStore.RETAIN_STALE`. If you evict the object so that you can still use it, but then you abort or

commit the transaction with `ObjectStore.RETAIN_STALE`, this cancels the retain type specified for `evict()`.

Maximum Data Lengths

For the `byte[] data` parameter in the constructor, the `MAXIMUM_DATA_LENGTH` variable specifies the maximum length of the byte array:

```
public static final int MAXIMUM_DATA_LENGTH = 16383
```

For the `String message` parameter, this variable specifies the maximum length of the UTF-8 encoding of the string. If the `String` is ASCII, the compression is one to one.

Restriction on data Argument Content

The C++ interface to `ObjectStore` treats the `data` argument as a C or C++ string, so there cannot be embedded zeros. If your application is directly providing a data array, it must ensure that there are no zeros in the data array. If your application is providing the data in the form of a string, it works correctly because the UTF-8 encoding of strings never uses zero bytes. If you try to construct a notification with a `byte[] data array` and you include a zero, `ObjectStore` signals `IllegalArgumentException`.

Subscribing to Receive Notifications

A session uses the `Notification.subscribe()` method to register to receive notifications for the specified location or within the specified segment or database. The database that contains the referenced object, cluster, or segment must be open when a session calls the `subscribe()` method. The overloads of `subscribe()` follow:

- `public static void subscribe(Placement placement)`
- `public static void subscribe(Object location)`

The `placement` parameter can specify a database, segment, or cluster. The subscribing session receives a notification if the application sends a notification that refers to an object in the specified cluster, segment, or database.

The `location` parameter specifies a persistent object.

A session can subscribe to many locations, clusters, segments, and databases simultaneously. ObjectStore stores subscriptions in the ObjectStore server for as long as the corresponding database is open for the associated session.

Discarding Subscriptions

When a session closes a database, ObjectStore discards any subscriptions in that session to notifications related to that database.

Unsubscribing from Notifications

A session can unsubscribe from particular notifications, just as it can subscribe to them. The unsubscription is immediate. The `Notification.unsubscribe()` method has the following overloads that are parallel to those for `subscribe()`:

- `public static void unsubscribe(Placement placement)`
- `public static void unsubscribe(Object location)`

If you try to unsubscribe from a notification that you are not subscribed to, nothing happens.

ObjectStore processes notifications asynchronously. Consequently, a session might unsubscribe from a particular notification but still receive notifications for that location because they were already queued.

The only way to cancel a cluster, segment, or database subscription is to unsubscribe from that cluster, segment, or database. You cannot unsubscribe from a cluster, segment, or database by unsubscribing from the notifications about the objects in that cluster, segment, or database.

Sending Notifications

An application can use any of the following methods to send notifications:

- `public void notifyImmediate()`
- `public void notifyOnCommit()`
- `public static void notifyImmediate(Notification[] notifications)`
- `public static void notifyOnCommit(Notification[] notifications)`

You can send one notification or an array of notifications at a time. You can send notifications immediately or when a transaction commits. If you wait to send a notification until a transaction commits, you ensure that a subscriber does not receive the notification until any associated change is visible in the database.

The object referenced in the notification cannot be stale and cannot have been destroyed.

Modification of an object does not cause a notification to be sent. A notification is sent only when the application program explicitly uses the notification API to send one.

Retrieving Notifications

The ObjectStore server queues notifications in the cache manager that is associated with the subscribing session. The session retrieves notifications from the cache manager. Typically, a session dedicates a thread to retrieve notifications from the queue. To do this, a thread calls `Notification.receive()`. The method signature is

```
public static Notification receive(int timeout)
```

The *timeout* parameter specifies the number of milliseconds to wait for a notification. Specify `ObjectStore.WAIT_FOREVER` to instruct the thread to wait forever. A value of 0 instructs the thread to return if there are no notifications in the queue.

The method returns null if there are no notifications in the allotted time. When the *timeout* parameter is `ObjectStore.WAIT_FOREVER`, ObjectStore never returns null.

When you call `receive()`, it does not matter whether or not a transaction is in progress. However, the thread from which you call `receive()` must be associated with a session.

A thread that calls the `Notification.receive()` method allows you to avoid polling your application. This method returns as soon as a notification is available, and until then, it waits. No polling is involved. The application is awakened when a notification arrives.

Reading Notifications

To extract the contents of a notification, use the following `Notification` class methods:

- `public Object getLocation()` obtains the persistent object.
- `public int getKind()` obtains the value for the `kind` parameter.
- `public byte[] getData()` returns the byte array associated with the notification. If the application specified a string when it created the notification, this method still returns the byte array. The notification itself does not maintain information about whether it was created with the specification of a string or a byte array.
- `public String getMessage()` decodes the byte array and returns the string associated with the notification. It does not matter whether you specified a string or a byte array when you created the notification. `ObjectStore` signals `java.io.UTFDataFormatException` if the string is not encoded correctly in UTF-8 format.

Managing the Notification Process

Managing the notification process involves consideration of the following:

- Notification Queue
- Performance Considerations
- Network Service

Notification Queue

The cache manager maintains a queue of notifications for each session on a machine. To set the size of the queue, call

`Notification.setQueueSize(int queueSize)`. `ObjectStore` signals `NotificationException` if a session subscribes to a notification before calling this method. The new queue size must not exceed the value specified for the `MAXIMUM_NOTIFICATION_QUEUE_LENGTH`.

Nothing forces a session to retrieve or read notifications. A session can subscribe to notifications but never retrieve any.

To avoid resource exhaustion in the cache manager, the size of the notification queue for each client is fixed. If the cache manager receives a

notification and the queue is full, ObjectStore discards the notification. This is called an overflow. Overflows do not cause any exception to be signaled and do not cause the application, the cache manager, or the server to crash.

The cache manager keeps statistics on the notification queue that include

- Queue size
- Number of pending notifications
- Number of overflows

To obtain these statistics, you can call the following `Notification` class methods:

- `public static int getQueueSize()`
- `public static getPendingNotifications()`
- `public static int getQueueOverflows()`

The information returned by these methods applies only to the session in which the method is called.

You can also use the ObjectStore utility `oscmstat` to obtain these statistics. The `ossvrstat` utility displays statistics on the number of notifications received and sent by the server. You can find more information about these utilities in *Managing ObjectStore*.

Performance Considerations

All notifications and subscriptions on a database go to the ObjectStore server. The server routes notifications to subscribed sessions through the cache manager, which queues the notifications for sessions. Because the server acknowledges each notification, sending a notification requires a round-trip message to the server.

Retrieving notifications accesses only shared memory and is very fast. If a session does not retrieve its notifications, the cache manager can run out of queue space. This causes the cache manager to discard notifications.

Every call to `subscribe()`, `unsubscribe()`, `notifyImmediate()`, and `notifyOnCommit()` requires one round-trip message to the ObjectStore server.

If any commit-time notifications are queued during a transaction, there is an additional remote procedure call (RPC) to the ObjectStore server during the commit operation.

Notifications are stored and forwarded in the server, cache manager, and sometimes even in the receiving application. Therefore, delivery of notifications might not be particularly fast. Performance varies according to system load and the amount of notification processing. For example, delivery could range from milliseconds to several seconds.

As a general rule, if you plan your application to use notifications, you should not expect high throughput. Do not expect a client application to send or receive more than about 10 notifications per second. In other words, do not use ObjectStore notification to meet high-speed requirements. The notification facility has a fairly simple design with limited buffering capability.

Network Service

When an ObjectStore application uses notifications, it automatically establishes a second network connection to the cache manager daemon on the local host. The application uses this connection to receive, and acknowledge the receipt of, incoming notifications from the cache manager. (Outgoing notifications are sent to the server, not the cache manager.) For specific information about defaults, see *Managing ObjectStore*.

Chapter 12

Using the Java Dynamic Data (JDD) Classes

The Java Dynamic Data (JDD) classes provide the API for creating, storing, and accessing persistent data, based on type information (schema) that the application defines at run time. JDD is designed for applications that model dynamic data and require greater flexibility when defining and redefining types than is available with persistence-capable Java classes.

This chapter explains how JDD works and shows how to use the JDD classes to perform common database operations.

Contents	This chapter covers the following topics:	
	An Overview of JDD	291
	Basic JDD Tasks	294
	Relationships	299
	Improving Query Performance with Superindexes	306
	Mixing Java Objects with JDD	307
Note	To compile and run a Java program that uses the JDD classes, you must set the <code>CLASSPATH</code> environment variable to include the location of the JDD .jar file.	

An Overview of JDD

Three concepts are central to understanding JDD: *types*, *attributes*, and *entities*. These concepts are analogous to Java's *classes*, *fields*, and *objects*. Just as a Java class defines a set of fields that can be assigned values in an object

of the class, so a JDD type defines a set of attributes that can be assigned values in an entity of the type.

The difference between Java and JDD is that, in Java, the definition of classes and their fields must occur *before* run time — during program development. In JDD, on the other hand, the definition of types and their attributes occurs dynamically at run time. Furthermore, JDD enables you to add, change, or remove types and their attributes at run time — without having to run the postprocessor, perform schema evolution, restart the Java virtual machine, or even begin a new transaction. In fact, JDD makes it possible to create a new type, add attributes to the type, create entities of the type, assign attribute values in the entities, and then perform queries — all in the same transaction. See A Simple JDD Application on page 297 for an example program that uses JDD to perform these tasks within a single transaction.

The following sections provide more information about JDD types, attributes, and entities.

Types

A type is an object of the class `Type` that defines the attributes that can be accessed for entities of the type. Once you have created a type by invoking the `Type ()` constructor, you can add attributes and create entities of the type.

Types have multiple inheritance; each type can have multiple subtypes and multiple supertypes.

Each type also maintains an extent of its entities and all of its subtypes' entities. This feature allows you to access entities through their type, either directly or by querying the type. A query on a type executes through the type's extent, which includes the extent of each of its subtypes.

Attributes

An attribute is an object of the class `Attribute`. You create attribute objects when you define a type's attributes, and you assign attribute values in the entities of the type. The different subclasses of `Attribute` — for example, `IntAttribute` and `EntityAttribute` — determine the types of values that can be assigned to attributes. These subclasses provide typed versions of accessor methods — for example, `get ()` and `put ()` — for type safety.

Methods in the `Attribute` class and its subclasses enable you to do the following:

- Set default values.

- Add constraints, including (for numeric attributes) maximum and minimum values.
- Declare an attribute as transient.
- Make an attribute immutable. Once the value of an immutable attribute is assigned, it cannot be changed.
- Make an attribute required. A required attribute must be assigned a value for each entity of the type.

Relationships

You can also use attributes to represent bidirectional relationships. The `Relationship` interface and the classes that implement it enable you to model three types of relationships:

- One-to-one
- One-to-many
- Many-to-many

To create a relationship, you invoke the constructor for the appropriate relationship class. The constructor creates attributes in the types on both sides of the relationship. Each attribute references the other side of the relationship.

For example, the `OneToMany()` constructor could be used to create a one-to-many relationship between two types, `Department` and `Employee`. The constructor would create one attribute in the `Department` type and the other in the `Employee` type. The attribute for `Department` could contain references to all `Employee` entities who are members of a department, and the attribute for `Employee` could reference the `Department` entity.

JDD maintains referential integrity on both sides of the relationship. To continue with the preceding example, if you were to add or remove an `Employee` entity, JDD would automatically update the attribute in the appropriate `Department` entity.

For more information, see Relationships on page 299.

Entities

An entity is an object of the `Entity` class. You create an entity as belonging to a type, and assign values to an entity's attributes, which are defined by the type. Later, when you add new attributes to the type, you can assign them values in the existing entities. You can also change an entity's type.

To make an entity available to a query, you must add it to the extent of the entity's type — by calling the `addToExtent()` method on the type. You can get a list of all the entities in the extent of a type — exclusive of its subtypes — by invoking the `entities()` method on the type. You can also get a list of all entities in the extent of the type and its subtypes, by invoking the `extent()` method.

Basic JDD Tasks

This section describes how to use the JDD API to perform basic database operations with JDD, including

- Defining Types and Their Attributes
- Creating Entities of a Type
- Querying a Type

Most of the code examples shown in these sections are from the sample program listed at the end of this section; see *A Simple JDD Application* on page 297.

Defining Types and Their Attributes

Before storing any entities in a database, you must first create their types. To create a type, invoke the `Type()` constructor, as in the following example:

```
Type Car = new Type(db, "Car");
```

The `db` argument is the database that you previously opened or created. The `Car` argument is the name of the new type, which is also the object returned by the constructor. To make `Car` a subtype of an existing type (for example, `Vehicle`), you specify an object of the supertype as the first argument, as in the following example:

```
Type Car = new Type(Vehicle, db, "Car");
```

After creating a type, you can create its attributes. It is not necessary that you create all attributes that the type will ever require. In fact, the value of JDD is that it lets you add attributes on an ongoing basis, as dictated by the changing data model. Likewise, you can remove attributes whenever the need arises.

To create an attribute, you invoke the constructor for one of the subclasses of the `Attribute` class. The choice of the subclass depends on the type of values

you will be storing in the attribute; for example, to store floating-point values, invoke the `DoubleAttribute()` constructor.

The following code defines three attributes in the `Car` type — `make`, `color`, and `year`:

```
StringAttribute make = new StringAttribute(Car, "make");
StringAttribute color = new StringAttribute(Car, "color");
IntAttribute year = new IntAttribute(Car, "year");
```

Indexing

If you will be performing query operations on the type, you can optimize queries by adding indexes. As explained in *Types* on page 292, you perform queries on a type, which maintains an extent of entities. To add an index, call the `addIndex()` method on the type, specifying the name of the attribute to be used as the key. The attribute name (a string) must be preceded by the dollar-sign character (`$`). You can cascade through a hierarchy of attributes by preceding each name with the `$` character; for example, `$a.$b.$c`.

The following example adds indexes to the `Car` type, specifying three attributes as keys — `make`, `color`, and `year`:

```
Car.addIndex("$make");
Car.addIndex("$color");
Car.addIndex("$year");
```

If `Car` has subtypes, `addIndex()` will also add indexes to the subtypes. Later, if you add another subtype, JDD automatically adds an index to the subtype, using the same keys.

For most applications, the indexing provided by `addIndex()` is sufficient to optimize queries. However, if your application queries types that have many subtypes, you should consider adding superindexes; for more information, see *Improving Query Performance with Superindexes* on page 306.

After you have defined types and their attributes, you have provided the type information needed to create entities.

Creating Entities of a Type

Creating an accessible entity requires not just the construction of an entity, but also assigning values to its attributes and adding the entity to its type's extent. This last step makes the entity available to queries.

To create an entity, call the `create()` method on its type. This method invokes the `Entity()` constructor and returns an `Entity` object, as follows:

```
Entity car = null;
car = Car.create();
```

If you do not have a type object available for calling the `create()` method, you can call the static `findType()` method to get the type, as follows:

```
Type Car = Type.findType("Car", db);
```

Assigning attribute values

To assign attribute values in the entity, call the `put()` method on the attribute object, as follows:

```
make.put(car, carMake[i%4]);
color.put(car, carColor[i%3]);
year.put(car, 1980+i);
```

If you do not have the attribute object (for example, `make`), you can call `put()` on the entity (`car`), specifying the name of the attribute as an argument, as follows:

```
car.put("make", "Ford"); // allowed, but not recommended!
```

However, such calls can be expensive, especially when assigning values to the same attribute in a highly iterative loop. Calling an entity's `put()` method with the attribute name as an argument incurs the expense of a string lookup for each call.

Instead, call `put()` on the attribute object, specifying the entity object and the value to assign as arguments. If you do not have the attribute object, call the `findAttr()` method on the type to retrieve the attribute, as in the following example:

```
// make this call outside the access loop
StringAttribute make = (StringAttribute)Car.findAttr("make");

// make this call to assign the attribute value inside the loop
make.put(car, "Ford");
```

The same consideration applies when calling `get()` to retrieve an attribute value.

Adding an entity to the type's extent

The final stage in creating an entity is to add it to the extent of its type. Unless you add it to the extent, the entity is not put in the database. The following example adds a `car` entity to its type, `Car`:

```
Car.addToExtent(car);
```

At this point, you have populated the database and can perform queries.

Querying a Type

Queries are performed in JDD using the query mechanism provided by OSJI; for detailed information, see [Querying ObjectStore Utility Collections](#). The only difference between OSJI queries and JDD queries is that, when you

specify an attribute in a query string, you must precede the attribute name with the \$ character, just as you do when adding an index; see *Defining Types and Their Attributes* on page 294. You can cascade references through a hierarchy of attributes by preceding each attribute with the \$ character; for example, \$a.\$b.\$c.

To query the database, construct a `TypeQuery` object, as follows:

```
TypeQuery query = new TypeQuery
    (Car, "($make == \"Fiat\") && ($year > 1990)");
```

You can use the `iterator()` method on the `TypeQuery` object to get a Java `Iterator` object and then use methods on the `Iterator` object to access the entities returned by the query.

A Simple JDD Application

The following sample program exercises the basic features of a JDD application to store and retrieve (through a query) information about cars.

The program creates a `Car` type and gives it three attributes: `make`, `color`, and `year`. The program then creates several `Car` entities, assigns values to their attributes, and queries the `Car` type for a list of entities that match the selection criteria specified by the query.

The command lines for compiling and executing the program follow the program listing.

```
// Cars.java: use the JDD classes to store information about
// cars; all type information is created at run time
import com.odi.*;
import com.odi.util.*;
import java.util.*;
import com.odi.jdd.*;
import com.odi.jdd.rel.*;

public class Cars {

    static Database db;
    static Transaction tr;

    static public void main(String args[]) {
        Session session = null;
        try {
            session = Session.create(null, null);
            session.join();
            load();
        } finally { session.terminate(); }
    }
}
```

```

static void load() {
    String carMake[] =
        new String[] { "BMW", "Honda", "Edsel", "Fiat" };
    String carColor[]=new String[] { "red", "black", "green" };

    System.out.println("Creating database ...");
    try {
        Database.open("cars.db",
            ObjectStore.UPDATE).destroy();
    } catch (DatabaseNotFoundException e) { }

    db = Database.create("cars.db",
        ObjectStore.ALL_READ | ObjectStore.ALL_WRITE);

    tr = Transaction.begin(ObjectStore.UPDATE);

    System.out.println("Generating type information ...");

    // create the type of the entities
    Type Car = new Type(db, "Car");

    // create attributes
    StringAttribute make = new StringAttribute(Car, "make");
    StringAttribute color = new StringAttribute(Car, "color");
    IntAttribute year = new IntAttribute(Car, "year");

    // add indexes to the attributes
    Car.addIndex("$make");
    Car.addIndex("$color");
    Car.addIndex("$year");

    System.out.println("Creating Entities ...");
    Entity car = null;
    for (int i = 0; i < 20; i++) {
        // create a Car entity
        car = Car.create();

        // assign values to its attributes
        make.put(car, carMake[i%4]);
        color.put(car, carColor[i%3]);
        year.put(car, 1980+i);

        // add the entity to its type's extent
        Car.addToExtent(car);
    }

    System.out.println(
        "Querying for all Fiats older than 1990 ...");
    TypeQuery query = new TypeQuery
        (Car, "($make == \"Fiat\") && ($year > 1990)");
}

```

```

        Iterator iter = query.iterator(null);
        System.out.println("\n  Make\tColor\tYear\n");
        while (iter.hasNext()) {
            car = (Entity)iter.next();
            System.out.println("  " + make.get(car) +
                               "\t" + color.get(car) + "\t" + year.get(car));
        }
        tr.commit(); // done
    }
}

```

Here are the command lines to compile and execute the program, along with the output from the run:

```

C:\examples>javac Cars.java
C:\examples>java Cars
Creating database ...
Generating type information ...
Creating Entities ...
Querying for all Fiats older than 1990 ...

    Make   Color  Year
    Fiat   black  1999
    Fiat   green  1991
    Fiat   red    1995

```

Relationships

The `Relationship` class and its subtypes enable you to create and use bidirectional relationships as attributes. The relationship classes provide methods for adding related entities to the attributes as well as for accessing them.

The types of relationships and the classes that implement them are as follows:

- One-to-one, implemented by the `OneToOne` class
- One-to-many, implemented by the `OneToMany` class
- Many-to-many, implemented by the `ManyToMany` class

JDD provides automatic maintenance on relationships whenever a member object is added or removed.

The following sections describe the relationship types and how to use them in a JDD application.

One-to-One Relationships

The one-to-one relationship is implemented by the `OneToOne` class. The constructor for this class creates attributes on both sides of the relationship, each of which is represented by an entity reference to the other side. The signature for the constructor is

```
public OneToOne(Type OneSide,
    java.lang.String OneSideAttribute,
    Type OtherSide,
    java.lang.String OtherSideAttribute)
```

You can use the constructed `OneToOne` object as an attribute object. For example, to assign related entities to the relationship, you would invoke the `put()` method on the `OneToOne` object. The signature of `put()` is

```
void put(Entity OneSideEntity, Entity OtherSideEntity)
```

To access the entity in the attribute, call the `get()` method on the attribute. Invoking the `OneSideAttribute.get()` method returns `OtherSideEntity`, and invoking the `OtherSideAttribute.get()` method returns `OneSideEntity`.

One-to-Many Relationships

The one-to-many relationship is implemented by the abstract class `OneToMany`. A one-to-many relationship has an owner (the “one” side) and members (the “many” side).

To create a one-to-many relationship, use the constructor for one of the subclasses, `LinkedOneToMany` or `IndexedOneToMany`. Their signatures are as follows:

```
public LinkedOneToMany(Type ownerType,
    java.lang.String ownerAttr,
    Type memberType,
    java.lang.String memberAttr,
    boolean doublyLinked)

public IndexedOneToMany(Type ownerType,
    java.lang.String ownerAttr,
    Type memberType,
    java.lang.String memberAttr,
    java.lang.String primaryIndexPath)
```

Both constructors create attributes for the types `ownerType` and `memberType`. The attribute on the owner side is a set of entity references to `memberType` entities, and the attribute on the member side is an entity reference to an

`ownerType` entity. For both constructors, the constructed object can be used as an `ownerAttr` object.

The constructor signatures for `LinkedOneToMany` and `IndexedOneToMany` differ in their last argument, as follows:

- The `doublyLinked` argument to `LinkedOneToMany()` specifies the type of list to use for the member entity references. If the `doublyLinked` argument is `true`, the list is doubly linked; otherwise, it is singly linked. A doubly linked list is more efficient when removing elements.
- The `IndexedOneToMany()` constructor implements the `ownerAttr` attribute as an `OSTreeSet`. It is for use with a very large set that you expect to query. The `primaryIndexPath` argument specifies the primary index; if this argument is omitted, no primary index is created. For more information about `OSTreeSet`, see Description of `OSTreeSet` on page 163.

The sample program listed in Relationship Example on page 303 uses the `LinkedOneToMany()` constructor to implement a one-to-many relationship between `Book` entities and `Borrower` entities, as follows:

```
LinkedOneToMany books = new LinkedOneToMany
    (Borrower, "books", Book, "BorrowedBy", true);
```

In the next line, the constructed object (`books`) is used to call the `add()` method, which adds the `Book` and `Borrower` entities to the relationship:

```
books.add(borrower[i%2], book);
```

The `add()` method returns `false` if the entities were already related in this relationship; otherwise, it returns `true` and adds them to the relationship.

The next line retrieves the borrower of a book, by calling the `get()` method on `book`, specifying the name of the attribute (`BorrowedBy`) as an argument:

```
Entity person = (Entity)book.get("BorrowedBy");
```

The next line makes two calls. The first call is to invoke the `get()` method on `books` (an attribute of the `Borrower` type), which returns a `ToManySet` object. This object is the set of `Book` objects related to `person`. The `iterator()` method is then invoked on the `ToManySet` object and returns an `Iterator` object that can be used to access the list of borrowed books:

```
Iterator iter = books.get(person).iterator();
```

Many-to-Many Relationships

The many-to-many relationship is represented by the abstract `ManyToMany` class. An object of this class maintains sets of entities on both sides of the

relationship. Whenever an element is added to the set that represents one side of the relationship, JDD maintains the inverse direction as well. Similar relationship maintenance is performed when removing elements from a set.

To create a many-to-many relationship, call the `ManyToMany()` constructor. Its signature is as follows:

```
public ManyToMany(Type ownerType,
    java.lang.String ownerName,
    Type memberType,
    java.lang.String memberName)
```

To add related entities to the relationship, call the `add()` method on the `ManyToMany` object. Its signature is as follows:

```
public boolean add(Entity ownerEntity, Entity memberEntity)
```

The `add()` method returns `false` if `ownerEntity` and `memberEntity` were previously linked in this relationship; otherwise, it adds entity references to the attributes and returns `true`.

Linked Objects and Many-to-Many Relationships

Some applications require a many-to-many relationship that can also represent the linking of two entities in a *link object*. The `ManyToManyWithObject` class can be used to implement both the relationship as well as link objects.

Consider, for example, an application that tracks students and the courses in which they are enrolled. The application creates two types, `Student` and `Course`. It implements a many-to-many relationship so that a `Course` entity can reference all the students taking the course, and a `Student` entity can reference all courses for which a student is enrolled. Furthermore, the application uses a link object to represent each student's enrollment in a course. The link object has a type (for example, `Enrollment`), enabling it to have attributes (for example, `grade`).

The `ManyToManyWithObject()` constructor has the following signature:

```
public ManyToManyWithObject(Type ownerType,
    java.lang.String ownerName,
    Type memberType,
    java.lang.String memberName,
    java.lang.String linkObjectType)
```

The `linkObjectType` argument is the type of the link object. If a type with that name already exists, JDD uses the existing type; otherwise, it creates a new type.

The constructor implements the relationship as two maps, one on each side of the relationship. The keys of each map are the related entities, and the values are the link objects. To illustrate with the previous example, each `Student` entity can have a map keyed by `Course` entities with values of `Enrollment` objects. And each `Course` entity can have a map keyed by `Student` entities, whose values are also `Enrollment` objects.

Accessor methods can be used to do the following:

- Retrieve the map of link objects.
- Get the set of entities referenced in the attribute.
- Add related entities to the relationship and create a link object.
- Unlink or remove entities from the map and delete the link object.

Relationship Example

The example program listed in this section is a simple lending- library application that keeps a record of borrowers and the books each has borrowed. It creates two types, `Book` and `Borrower`, and uses the `LinkedOneToMany` class to implement a one-to-many relationship that represents the relationship between borrower and book. A borrower can have many books, but a book can have only one borrower.

The `Book` and `Borrower` types each have two attributes. The `Book` type has these attributes:

- `title`, the title of a `Book` entity
- `BorrowedBy`, a reference to a `Borrower` entity

The `Borrower` type has these attributes:

- `name`, the name of a `Borrower` entity
- `books`, the list of the `Book` entities loaned out to `Borrower`

The `title` and `name` attributes are created by invoking the `StringAttribute()` constructor. The `BorrowedBy` and `books` attributes are created by invoking the `LinkedOneToMany()` constructor.

Each of the `Book` entities is created in a `for` loop. After an entity is created, the book it represents is recorded as having been loaned out to a borrower by adding the related `Book` and `Borrower` entities to the one-to-many relationship.

The command lines for compiling and executing the program are shown following the source code listing.

```
// Library.java: use JDD classes to store and retrieve
// information for a lending library -- who has borrowed which
// books. The one-to-many relationship between borrower and
// books is implemented by JDD's LinkedOneToMany class.
import com.odi.*;
import com.odi.util.*;
import java.util.*;
import com.odi.jdd.*;
import com.odi.jdd.rel.*;

public class Library {

    static Database db;
    static Transaction tr;

    static public void main(String args[]) {
        Session session = null;
        try {
            session = Session.create(null, null);
            session.join();
            loan();
        } finally { session.terminate(); }
    }

    static void loan() {
        Entity borrower [] = new Entity[2]; // borrowers

        String titleStr[] = new String[] { // titles of books
            "Life of Johnson", "Tale of a Tub", "Rambler",
            "Beggar's Opera", "Memoirs of Martinus Scriblerus",
            "Reflections", "Seventeen Thirty-Eight"
        };

        // create database
        try {
            Database.open
                ("library.db", ObjectStore.UPDATE).destroy();
        } catch (DatabaseNotFoundException e) { }

        db = Database.create("library.db",
            ObjectStore.ALL_READ | ObjectStore.ALL_WRITE);

        tr = Transaction.begin(ObjectStore.UPDATE);

        // create Book type with title attribute
        Type Book = new Type(db, "Book");
        StringAttribute title = new StringAttribute(Book, "title");
```



```

// create Borrower type with name attribute
Type Borrower = new Type(db, "Borrower");
StringAttribute name =
    new StringAttribute(Borrower, "name");

// create a one-to-many relationship between Borrower and
// Book
LinkedOneToMany books = new LinkedOneToMany
    (Borrower, "books", Book, "BorrowedBy", true);

// create two Borrower entities
borrower[0] = Borrower.create();
borrower[0].put(name, "Terry");
Borrower.addToExtent(borrower[0]);
borrower[1] = Borrower.create();
borrower[1].put(name, "Alice");
Borrower.addToExtent(borrower[1]);

// create Book entities
Entity book = null;
for (int i = 0; i < 7; i++) {
    book = Book.create();
    title.put(book, titleStr[i]);
    Book.addToExtent(book);

    // lend them out to the borrowers
    books.add(borrower[i%2], book);
}

// Who borrowed the last book?
Entity person = (Entity)book.get("BorrowedBy");
System.out.println("\n" + title.get(book) +
    "\n" is out on loan to " + person.get(name));

// get the list of books loaned to this borrower
Iterator iter = books.get(person).iterator();
System.out.println(person.get(name) +
    " borrowed these books:");
while (iter.hasNext())
    System.out.println("\t" +
((Entity)iter.next()).get(title));

tr.commit(); // done
}
}

```

Here are the command lines to compile and execute the program, along with the output from the run:

```

C:\examples>javac Library.java
C:\examples>java Library

```

```
"Seventeen Thirty-Eight" is out on loan to Terry
Terry borrowed these books:
  Life of Johnson
  Rambler
  Memoirs of Martinus Scriblerus
  Seventeen Thirty-Eight
```

Improving Query Performance with Superindexes

A superindex is a special type of index for optimizing queries on types that have many subtypes. The following sections provide background information to explain how JDD implements default indexing; when default indexing may not be sufficient to optimize queries; and how superindexes can improve the performance of queries.

Note Applications that do not query types with many subtypes will probably not benefit from superindexing and should use the default indexing as described in *Defining Types and Their Attributes* on page 294.

Queries and Default Indexing

When an application queries a type, by default the query is recursive: separate query operations execute on the type and on any of its subtypes. Likewise, when you call `addIndex()` on a type, the index is recursively added to the type and to all its subtypes. These indexes are automatically maintained and updated by JDD when you add or remove a subtype, or when you change the value of an attribute.

This default level of indexing is sufficient for most query operations. However, if you query a parent type that has a large and intricate hierarchy of subtypes, indexing can add to the overhead of the query. To reduce the time it takes to query such types, JDD provides the superindex.

Superindexing

When you call `addSuperIndex()` on a type, JDD adds a single superindex to the type, which indexes the parent type and all its subtypes. When you query the parent type, only one query operation occurs — not the recursive queries that occur when you query a type with default indexing. As with default indexing, JDD automatically updates the superindex if you add or remove a subtype or if you change the value of an attribute.

There are several disadvantages to using a superindex that you should consider when deciding which type of index to add:

- A superindex can benefit query performance only when added to types having many subtypes. If your application queries types with few or no subtypes, adding a superindex will not significantly improve query performance.
- Invoking the `addSuperIndex()` method creates a single superindex, which is added to the type object to which the method is applied. If you begin the query with a subtype, you do not get the benefit of *any* indexing — unless you have explicitly called `addSuperIndex()` on the subtype.
- When you change an attribute value for a type that has a superindex, JDD automatically updates any superindexes that have been added to the subtypes. If you have added many superindexes to many subtypes, index maintenance can become time consuming.
- Superindexes may not be added to types that have subtypes in different segments of the same database or in different databases. Once a superindex has been added to a type, no new subtypes may be added that are in a different database or segment.

Mixing Java Objects with JDD

JDD enables you to mix persistent Java objects and JDD entities in the same application. JDD classes are extendable. For example, you can extend the `Entity` class to have native Java fields and methods. You can similarly extend the `Type` and `Attribute` classes.

When you extend JDD classes to include native fields, however, you lose some of the flexibility that JDD provides. The reason is that, to make the extended classes persistence-capable, you must use the `ObjectStore` API to do the following:

- Run the `osjcfp` postprocessor to make the classes persistence capable.
- Perform schema evolution if you change any of the persistence-capable classes.

If you only have to run the postprocessor once and never have to perform schema evolution, these tasks may not be an issue. But if your application models data that is constantly changing, you will probably want to rely as much as possible on JDD for persistent storage.

This section is organized as follows:

- When You Must Use the ObjectStore API on page 308
- Pros and Cons of Mixing the ObjectStore API with JDD on page 308
- Using Extended JDD Classes on page 308

When You Must Use the ObjectStore API

You must use the ObjectStore API to perform the following operations:

- Creating and managing sessions; for more information, see *How Sessions Keep Threads Organized*.
- Opening and closing a database; for more information, see *Chapter 4, Managing Databases*.
- Starting and committing a transaction; for more information, see *Chapter 5, Working with Transactions*.

Pros and Cons of Mixing the ObjectStore API with JDD

Among the reasons for using the ObjectStore API to include persistence-capable classes in a JDD application are the following:

- Native fields of Java objects take up much less space in the database than JDD attributes.
- Accessing and modifying native fields is significantly faster.

If either of these factors weighs heavily with your application, you might consider extending JDD classes to define native fields and accessor methods.

Fields vs. attributes

When deciding whether to store data as a native field or an attribute, you should consider whether you will have to change the field at a later date. Deleting or renaming a field or changing its type requires recompiling source code, running the postprocessor, and possibly performing schema evolution. Changing an attribute, on the other hand, has no such risks.

Access time

Another factor to consider is access time: accessing a native field is significantly faster than accessing an attribute. If you need to access the same data in a highly iterative, time-critical loop, you should consider storing it as a native field.

Using Extended JDD Classes

The most likely candidate for extending is the `Entity` class. After compiling and postprocessing an application that uses the extended entity, you use the

application to create a type for the entity, give the type a set of attributes, and create entities of the type — using the same procedure described in Defining Types and Their Attributes on page 294 and Creating Entities of a Type on page 295.

Creating attribute objects for native fields

If you want to use any of the JDD accessor methods on a native field, you must first create an attribute object for the field, as described in Creating Entities of a Type on page 295. Note, however, that using a JDD method to access a native field is significantly slower than using a native method. When you use a JDD accessor method on an attribute object associated with a native field, JDD uses Java reflection to access the field.

Querying native fields

JDD allows you to mix native fields and attributes in the same query string. The query treats the field reference differently, depending on whether or not you precede the name with the `$` character:

- If the name begins with the `$` character (for example, `$age`), JDD interprets it as the name of an attribute object. If the object is associated with a native field, the query uses Java reflection to access the field.
- If the name does *not* begin with the `$` character (for example, `age`), JDD interprets it as the name of a native field. The query accesses the field directly, just as an ObjectStore query would.

Indexing native fields

Creating attribute objects for native fields also allows you to add indexes to the fields. JDD performs the same index maintenance on native fields as on attributes, *except* if either of the following is true:

- The ObjectStore API was used to add the index.
- A field was updated directly instead of through a JDD accessor method.

If either condition is true, JDD does not update the index.

Example

The following example uses a persistence-capable class, `UserEntity`, which extends JDD's `Entity` class and defines a native integer field called `age`:

```
// create the type User in the database represented by db
Type User = new Type(db, "User");

// create an entity of the type User; UserEntity extends Entity
UserEntity user = new UserEntity(User);

// assign to the age field
user.age = 32; // native Java field
```

To use a JDD method to access `age`, you first create an attribute object for it, as follows:

```
// create an attribute object for the age field, which is
// defined in the UserEntity class
IntAttribute ageAttr =
    (IntAttribute)myUser.findAttr("age");

// call JDD's put() method on the attribute object for the age
// field
ageAttr.put(user, 32);
```

Chapter 13

Miscellaneous Information

This chapter provides miscellaneous information about ObjectStore.

Contents

This chapter discusses the following topics:

Java-Supplied Persistence-Capable Classes	311
Description of Special Behavior of String Literals	315
Serializing Persistent Objects	317
Using Persistence-Capable Classes in a Transient Manner	319
Comparing Java and C++ Persistent Storage Layouts	319
Differences Between C++ and Java Interfaces to ObjectStore	321
Environment Variables	322

Java-Supplied Persistence-Capable Classes

Some Java-supplied classes are persistence capable. Others are not persistence capable and cannot be made persistence capable. A third category of classes can be made persistence capable, but there are important issues to consider when you do so.

Description of Java-Supplied Persistence-Capable Classes

The following Java classes are persistence capable:

- `java.lang.String`

The wrapper classes follow:

- `java.lang.Boolean`
- `java.lang.Byte`
- `java.lang.Character`
- `java.lang.Double`
- `java.lang.Float`
- `java.lang.Integer`
- `java.lang.Long`
- `java.lang.Short`
- Arrays of `Object`, of any of the primitive types (`boolean`, `byte`, `integer`, and so on), and of any persistence-capable type are all persistence capable. (You can allocate an array and initialize it later, just as you would with any other field.)

Identity

`ObjectStore` does not always preserve identity for objects that are instances of the Java wrapper or `String` classes. It is more efficient to store these objects as values rather than as objects. Because identity is not always preserved, programs that use object identity to compare wrapper class objects work differently when used with persistent objects. For example, the following method is incorrect:

```
boolean comparePersistIntegers(Integer x, Integer y) {  
    return (x == y);  
}
```

Instead, it should be written as

```
boolean comparePersistIntegers(Integer x, Integer y) {  
    return x.equals(y);  
}
```

Additional information about object identity is in [About Object Identity](#) on page 119.

Virtual memory overhead

When `ObjectStore` makes them persistent, `String` types, primitive wrapper types, and arrays have more run-time virtual memory overhead than types that implement `IPersistent`. This is because `ObjectStore` must create entries for these types in two hash tables. `com.odi.IPersistent` requires an entry in a single hash table because certain information is stored in fields in the object.

Persistent and persistence capable

In your program, some wrapper objects or strings might be persistent and some might be transient, though both are persistence capable:

- If the application explicitly calls `ObjectStore.migrate()` on a wrapper object or string and then stores it in a `com.odi.coll` collection, `com.odi.util.OSTreeSetxxx`, or as a value in an `com.odi.util.OSTreeMapxxx`, the wrapper object or string becomes persistent.
- If the wrapper object or string is only reachable through transitive persistence, it does not become persistent when the transaction is committed. Instead, `ObjectStore` stores the object as an immediate value. This means that `ObjectStore` does not store the object in any of its internal hash tables and does not store the object as a separate value in the database. Instead, `ObjectStore` stores the object in the location of the reference to the object. The reference completely describes the object.
- Some strings become immediate if you set the `com.odi.useImmediateStrings` property to `true` when a session is created and the string is eight characters or less. These strings are stored as values of fields in persistent objects and are not persistent unless you explicitly call `ObjectStore.migrate` to make them persistent.

Any routine that requires a persistent object, as opposed to a persistence-capable object, notices the distinction between persistent and *persistence capable but transient*. For example, if an application calls `Segment.of()` on an `Integer` object, the return value might be a segment in a database or `ObjectStore` might signal `ObjectNotPersistentException`. You cannot always predict what the return value will be because an `Integer`-valued field in a persistent object can contain either a persistent or a transient value.

Unicode strings

`ObjectStore` stores Unicode strings. You can specify any Java string with Unicode characters in it and `ObjectStore` can store it persistently and retrieve it correctly. `ObjectStore` uses UTF-8 (Unicode Transformation Format) encoding/compression to store regular English strings compactly. Sun's Java implementation uses the same mechanism.

Immediate strings

The identity of immediate strings is not maintained.

Can Other Java-Supplied Classes Be Persistence Capable?

Many Java system classes cannot be persistence capable. There are other Java system classes that you can make persistence capable but you must consider some issues when you do so. In some situations, you can subclass the Java system class and make the subclass persistence capable. Of course, this would not work for final classes.

Primitive types	<p>You cannot store an object of a primitive type, such as an <code>int</code>, directly in a database as a discrete object. To store an object of a primitive type in a database you can</p> <ul style="list-style-type: none">• Place it in a wrapper object, such as an <code>Integer</code>• Define it as a field in a persistence-capable class <p>For example, you cannot make <code>byte</code> persistence capable because by itself a <code>byte</code> is not an <code>Object</code>, but you can make <code>byte[]</code> persistence capable because it is an <code>Object</code>.</p>
Native methods	<p>Classes that use native methods cannot be made persistence capable by the postprocessor because the postprocessor cannot annotate the native methods the way it can annotate Java code. What this means is that if a class has native methods and you postprocess the class, <code>ObjectStore</code> cannot guarantee that everything will work properly.</p> <p>You might choose to postprocess your code, then add native code. If you do this, you must ensure that any persistent objects that your native code references are properly fetched from the database before the native method is called. Be careful, however, if a native method changes the value of an indexed instance variable. This does not work properly because the index is not updated.</p>
Classes that hold state	<p>Other system classes do not make sense as persistent objects because they hold state that is inherently tied to the process, such as open file channels or Java threads.</p>
Post-processing	<p>For other classes, such as <code>java.lang.Stringbuffer</code>, the previous obstacles might not apply. If you postprocess the <code>.class</code> file for <code>java.lang.StringBuffer</code> and specify the <code>-modifyjava</code> option, the postprocessor produces a persistence-capable <code>StringBuffer</code> class:</p> <pre>osjcfp -dest \osjcfpout -modifyjava java.lang.StringBuffer</pre> <p>Then you must put the new <code>.class</code> file in your <code>CLASSPATH</code> variable ahead of the standard Java <code>.class</code> file. All subsequent use of the <code>StringBuffer</code> class in this environment would use the persistence-capable version.</p>
Performance drawbacks	<p>There are, however, some drawbacks to doing this. Some slowdown of some or all the methods will occur, because the postprocessor must add new instructions to check whether the object needs to be brought in from the database or needs to be marked as modified.</p> <p>The amount of slowdown is hard to determine. It depends on the details of the method. Even parts of your program that never handle persistence are</p>

	<p>affected by these extra instructions. This also applies to indirect uses of the class, for example, if <code>StringBuffer</code> is used heavily in a Java library that you are using, such as a user interface or network library.</p>
Library version problems	<p>There can also be problems with Java library version skew. If you postprocess <code>java.lang.StringBuffer</code> from version 1.1 of the Java Virtual Machine, then your user uses your program with version 1.1.2, and <code>StringBuffer</code> has changed in some way between 1.1 and 1.1.2, your user will see the 1.1 version (persistence capable) everywhere in the entire Java environment. If your user was depending, directly or indirectly, on the new 1.1.2 version of <code>StringBuffer</code>, something might not work properly.</p>
Renaming the class	<p>You might need to rename the newly created persistence-capable version so that the non-persistence-capable version is still available to the other Java system classes. To do this, specify the <code>-translatepackage</code> option when you run the postprocessor. See Putting Processed Classes in a New Package on page 221.</p> <p>This avoids the problem and is generally safer. However, you might need the persistence-capable class to have the original class name. For example, suppose you have a library with a method that takes an argument of type <code>java.lang.Stringbuffer</code>. You want to pass in a persistence-capable object. You cannot rename the class because the argument type would not match.</p>
<code>java.util.Hashtable</code>	<p><code>ObjectStore</code> itself uses <code>java.util.Hashtable</code>. Consequently, invoking Java or using <code>ObjectStore</code> with a persistence-capable version of <code>java.util.Hashtable</code> that is available in your <code>CLASSPATH</code> is likely to cause trouble, such as infinite loops. A better approach is to substitute the <code>ObjectStore</code>-supplied class <code>com.odi.util.OSHashtable</code>.</p>

Description of Special Behavior of String Literals

There are special considerations when you make `String` literals persistent.

When a Java program refers to a `String` literal by using quotation marks to name a string, Java treats the resulting `String` as a constant value. Multiple calls to a method with the `String` literal operate on the same `String` object.

The `com.odi.stringPoolSize` initialization property allows you to control the way that `ObjectStore` causes `Strings`, other than literals with the same

contents, to be represented by a single, shared instance in the database in certain circumstances. See *Description of com.odi.stringPoolSize* on page 45.

Example of String Behavior

Consider the following example:

```
Object string() {String result = "string"; return result;}
Object intArray() {int[] result = { 1, 2, 3 }; return result;}
boolean stringsTheSame() {return string() == string();}
boolean intArraysTheSame() {return intArray() == intArray();}
```

The `stringsTheSame()` method always returns `true` because every call to `string()` returns the same `String` object. The `intArraysTheSame()` method always returns `false` because each call to `intArray()` constructs a new `int[]` object.

Destroying Strings

By default, `String` objects that become persistent during a transaction revert to being transient at the end of the transaction. Persistent objects usually are made stale at the end of a transaction. Unlike objects that implement `IPersistent`, when a `String` is made stale, it becomes transient.

When you destroy a `String`, in the transaction in which the destroy operation occurs, `ObjectStore` keeps track of the fact that the object was destroyed. An attempt to use a destroyed `String` literal causes `ObjectStore` to signal `ObjectNotFoundException`. The solution is to copy the `String` before you destroy it.

You should not destroy a `String` in a database unless you know that no other object in the database refers to that `String`. A safe, though possibly inefficient, way to handle this is to use

```
new String(String)
```

to force a new identity to each `String` that might be referenced. Also, you must disable the `String` pool by specifying `0` for the value of the `com.odi.stringPoolSize` initialization property. This ensures that you can safely destroy the old `String` instance.

It is usually best to avoid destroying strings or objects altogether and let the persistent garbage collector take care of destroying such unreachable objects. The persistent GC can typically destroy and reclaim such objects very efficiently, because it can batch such operations and cluster them effectively. If you set up the GC to run when the system is lightly loaded, you can defer the overhead of the destroy operations to a time when your system would

otherwise be idle, thus getting greater real throughput from your application when you really need it.

Immediate strings

You do not need to destroy immediate strings. Destroying an immediate string has no effect. Immediate strings are not persistent objects. They are transient values that are not stored in the database.

Serializing Persistent Objects

You can serialize many classes that implement `com.odi.IPersistent`. For this to work, the definition of your persistence-capable class must implement the `java.io.Serializable` interface. The classes you can serialize include

- `com.odi.util.OMHashtable`
- `com.odi.util.OMTreeMap`
- `com.odi.util.OMVector`
- `com.odi.util.OMTreeSet`

During serialization, none of the transient fields in the `IPersistent` implementation needs to be written out. When you deserialize instances of the `com.odi.util.OMTreeMap` and `com.odi.util.OMTreeSet` classes, you can specify their placement by using the `ObjectStore.setPlacementForSerialization()` method.

Before serializing an object, an application must always invoke `ObjectStore.deepFetch()` on the object to be serialized. The `deepFetch()` method ensures that the contents of all components of the object are accessible. This must be the case for an application to serialize an object.

Why you use deepFetch()

In an `ObjectStore` application, the first time you read or modify an object, `ObjectStore` makes the contents of the object available. The contents do not have to be available before you start the operation. You need not add Java code to make the contents available. When an `ObjectStore` program follows a reference from a source object to a target object, the contents of the target object are automatically available. This happens because the postprocessor recognizes the Java byte-code instructions that follow references and it inserts the code that fetches the object contents.

Serialization works differently. It follows references from one Java object to another without using Java byte codes. Serialization does not perform the automatic fetches the way that `ObjectStore` does. Consequently, before you initiate serialization of an object, its contents and the contents of all its

components must already be available. The `ObjectStore.deepFetch()` operation does this for you.

Limitation You cannot serialize Java peer objects. Consequently, you cannot serialize `ObjectStore` collection objects.

Example When an application serializes and deserializes a persistent object with the default serialization methods, `ObjectStore` effectively creates a transient copy of the object and its components. Following is code that provides an example of serializing and deserializing persistent objects. In this example, `list2` is a transient copy of the persistent list.

```
public
class SerializationExample {

public static void main(String argv[])
    throws java.lang.ClassNotFoundException,
    java.io.IOException,
    java.io.FileNotFoundException {
    String dbName = argv[0];
    Session.createGlobal(null, null);
    /* Create a database with a list in it. */
    Database db = Database.create(dbname,
        ObjectStore.ALL_READ | ObjectStore.ALL_WRITE);

    Transaction tr = Transaction.begin(ObjectStore.UPDATE);
    List curr = new List("1", null);
    db.createRoot("list", curr);
    for (int i=2; i < 5; i++) {
        curr.next = new List(""+i, null);
        curr = curr.next;
    }
    tr.commit();

    /* Illustrate use of serialization in this example. */
    tr = Transaction.begin(ObjectStore.UPDATE);
    List head = (List)db.getRoot("list");

    /* Fetch the entire list prior to serializing it. */
    ObjectStore.deepFetch(head);

    FileOutputStream f = new FileOutputStream("tmp");
    ObjectOutputStream os = new ObjectOutputStream(f);
    os.writeObject(head);

    FileInputStream in = new FileInputStream("tmp");
    ObjectInputStream is = new ObjectInputStream(in);

    /* list2 is effectively a copy of the list denoted by head. */
    List list2 = (List)is.readObject();
    ...

    tr.commit();
}
```

```

    }
}

public class List implements java.io.Serializable {
    public Object value;
    public List next;

    List(Object value, List next) {
        this.value = value;
        this.next = next;
    }

    ...
}

```

Using Persistence-Capable Classes in a Transient Manner

The `stublib.jar` file contains stubs of `ObjectStore` classes that allow user-defined persistence-capable classes to be used in a purely transient manner. The annotations in the persistence-capable classes make calls to the various `ObjectStore` stub routines in `stublib.jar`.

The `stublib.jar` file provides a stripped-down version of the `ObjectStore` API. This allows better performance and a smaller footprint than the complete `.jar` file.

For example, you might want to use `stublib.jar` for the client in an RMI or CORBA application. The client might use persistence-capable classes, which make references to various `ObjectStore` methods, but the client never directly accesses a `ObjectStore` database. In this situation, the stub routines in `stublib.jar` satisfy the requirements of the Java VM's linker.

To use `stublib.jar`, put it in your `CLASSPATH` instead of `osji.jar`.

If your application uses any classes in `com.odi.util`, you must use `osji.jar`. You cannot use `stublib.jar` because the stub definitions are not sufficient for the `com.odi.util` classes.

Comparing Java and C++ Persistent

Storage Layouts

There are differences between databases created by the Java and C++ interfaces to ObjectStore. These differences result in restrictions on the use of databases by both the C++ and Java interfaces.

Java databases

Databases created by the Java interface can be used by C++ programs, but the representation of Java primary objects is different from regular C++ objects. Because of these differences, accessing the contents of Java primary objects from C++ programs is not currently supported. C++ programs can store, access, and modify C++ objects in databases created or modified by the Java interface.

C++ databases

Databases created by the C++ interface can be read or modified by the Java interface. C++ objects can be manipulated as described in *Developing Java Applications That Access C++*. In addition, the Java interface can store Java primary objects in databases created by the C++ interface.

Java primary objects

Databases that hold Java primary objects have a segment that contains information used by the Java interface to describe the schema for the classes of the Java objects stored in the database.

All Java primary objects contain a 4-byte object header followed by data for the object fields. Java arrays are represented by a 12-byte header object and a separate C++ array that contains the array contents.

Java primitive fields

Java primitive fields are represented the same way as similarly sized C++ primitives. The sizes of Java primitives when they are stored in instance fields or arrays are shown in the following table.

<i>Primitive</i>	<i>Primitive Size in Bytes</i>
boolean	1
byte	1
short	2
char	2
int	4
float	4
double	8
long	8

Java object reference fields

Fields of Java primary objects that contain object references are represented by an 8-byte or 12-byte data structure rather than the 4-byte pointer usually used by C++ objects. The larger data structure allows the Java interface to provide more features for object references.

The 8-byte data structure is used for Java object reference fields that cannot contain arrays, `Strings`, `Doubles`, or `Longs`.

The 12-byte data structure is used for Java object reference fields that contain `Objects`, `Strings`, `Doubles`, `Longs`, and arrays.

Differences Between C++ and Java Interfaces to ObjectStore

Following are some differences between the Java and C++ interfaces to `ObjectStore`.

Timing of the Write Lock Acquisition

In the C++ interface to `ObjectStore`, as soon as you modify an object, you set or try to set the `ObjectStore` write lock for the page that the object is on. But in the Java interface to `ObjectStore`, this might or might not happen, depending on the lazy write locking flag. The default is that it does not happen, and the write locking is deferred until later. Thus, if you have two sessions (in two VMs) that are accessing the same data and they are not both just reading, different timing of the write lock acquisition can cause the behavior to be different in the Java interface than it is in the C++ interface.

Opening the Same Database Multiple Times

In the Java interface to `ObjectStore`, each subsequent opening of a database after the initial open operation returns the same database object. For example:

```
db1 = Database.open("foo", ObjectStore.UPDATE);
db2 = Database.open("foo", ObjectStore.UPDATE);
```

In the Java interface to `ObjectStore`, the expression `db1 == db2` returns `true`. They refer to the same database object. Consequently, a call to `db1.close()` or `db2.close()` closes the same database. No matter the number of times you open a database, a single call to the `close()` method closes the database.

This is different in the C++ interface to ObjectStore. In that interface, for example, if you call `open()` four times and `close()` three times all on the same database, the database is still open.

Environment Variables

ObjectStore includes the following environment variables:

- `OS_JAVA_VM` specifies the command for running the Java virtual machine, when it is set. The default is that this variable is not set. The Windows tool batch files use the value of this variable when it is set.
- `OSJCFPJAV` specifies the name of the Java executable you want the postprocessor to use. The default is `java`. If this variable is not set, the postprocessor uses the first Java executable that it finds in your `PATH` environment variable. If you want the postprocessor to use another Java executable, set the `OSJCFPJAV` environment variable to the name of the Java executable you want the postprocessor to use.

If the postprocessor cannot find a Java executable, it signals a `Bad command or file name` error message.

Chapter 14

Tools Reference

This chapter provides reference information for the following tools:

osjcfp: Running the Postprocessor	323
osjgen: Generating Peer Classes	333
osjcheckdb: Checking References in a Database	338
osjshowdb: Displaying Information About a Database	339
osjuphsh: Upgrading String Hash Codes in Databases	340
osjversion: Obtaining ObjectStore Version Information	341

osjcfp: Running the Postprocessor

To make classes persistence capable, compile the source files, then run the postprocessor on the resulting class files. You must run the postprocessor on all class files in a batch at the same time. The postprocessor can accept a command line that intersperses file names, options, and input file specifications. Complete information about the postprocessor is in Chapter 8, *Generating Persistence-Capable Classes Automatically*, on page 195.

Command Format

```
osjcfp -dest destination_dir file_name [file_name...][options].
```

Options

@input_file

Causes the contents of the named input file to replace this argument in the command line. The postprocessor does this before any other argument

processing. You can specify this option multiple times on one command line to include multiple files. You cannot nest this option. That is, the postprocessor does not expand this argument if it appears in an input file.

`-annotatefield qualified_field_name`

Instructs the postprocessor to insert `fetch()` or `dirty()` calls in annotated code whenever a method accesses a transient field in a persistent class. The `-annotatefield` option turns off the `-noannotatetransientfields` option.

`{ -cis | -classinfosuffix } suffix_string`

Specifies the suffix that the postprocessor adds to the name of the `ClassInfo` subclass that the postprocessor generates for each class that it makes persistence capable. By default, the suffix is `ClassInfo`. This is useful when you need to limit the number of characters in file names. For all batches in an application, you must specify the same suffix if you do not use the default.

`{ -cpath | -classpath } class_path`

Specifies the path by which to locate class files for postprocessing. If you specify this option, `ObjectStore` uses it in place of the `CLASSPATH` environment variable. The `-classpath` option does not affect the class path used to run the postprocessor.

`{ -sysclasspath | -sysclasspath } system_class_path`

Specifies the path by which to locate class files for Java system classes for postprocessing. If you specify this option, `ObjectStore` uses it in place of the value in the `sun.boot.class.path` system property when using the Sun JDK 1.2. When using `jview`, system classes cannot be found unless you specify their location using this option. The `-sysclasspath` option does not affect the class path used to run the postprocessor.

`{ -cc | -copyclass }`

Copies classes to the destination directory without annotating them. This option applies to class names, `.class` files, `.jar` files, and `.zip` files that you specify on the command line after the `-copyclass` option and before the next `-persistcapable` or `-persistaware` option or the end of the command line. This option is useful when you want nonpersistent classes or classes that have already been annotated to be in the same directory as persistence-capable or persistence-aware classes being created.

If you specify the `-persistaware` or `-persistcapable` option for any file for which you also specify the `-copyclass` option, the postprocessor ignores the `-copyclass` option for that file.

If you specify the `-translatepackage` option and the `copyclass` option, the postprocessor modifies the class to accommodate the new package name.

```
{ -d | -dest } destination_dir
```

This option is required. The postprocessor uses the directory you specify for *destination_dir* as the root for locating the annotated files. The postprocessor places each class file it operates on in the package-appropriate subdirectory of the destination directory, as though the destination directory were in your class path.

If the destination directory specification would cause the postprocessor to overwrite an original file and you did not specify the `-inplace` option, the postprocessor reports an error and terminates without producing any output.

```
{ -emlf | -embeddedmaxlengthfield }
```

Specifies that Java Strings for a given field will be represent as embedded Unicode arrays. Using embedded strings can significantly improve performance for those applications that heavily use small strings (less than 100 characters).

This option takes the following two arguments:

- First argument is the fully qualified name of a string field
- Second argument is the maximum size string that can be stored in that field.

Note that you cannot declare embedded string fields as indexable fields for Java peer collections.

```
{ -f | -force }
```

Forces the postprocessor to overwrite existing annotated `.class` and `ClassInfo` files.

```
-hashcode class_name
```

Causes the postprocessor to add a persistent `hashCode()` method to the specified class. You typically use this option with the `-nodefaulthashCode` option. If you specify this option for a class for

which you explicitly defined a `hashCode()` method, the postprocessor reports an error.

```
-includesummary inc_class_name
```

Instructs the postprocessor to include the specified summary in the new summary it creates. It does not matter whether the postprocessor is also annotating any classes. Replace *inc_class_name* with the name of a file generated by a previous execution of the postprocessor with the `-summary` option. You must know the name of the generated file. If you did not postprocess the library yourself, you must get the name of the generated class from the library vendor. When you specify the `-includesummary` option, you must also specify the `-summary` option.

```
{ -index | -indexablefield } field
```

Marks a field as indexable for a peer (`com.odi.coll`) collection. This option applies only to the field that immediately follows it. You must specify a fully qualified field, for example, `com.odi.demo.people.name`. You can specify this option multiple times. This option does not apply to utility (`com.odi.util`) collections.

This option is useful because it allows a query to run faster. The postprocessor does not actually add the index. You can add a persistent index with a call to the API at run time. When the index is present, queries that use the specified field are faster.

Suppose you declare a field to be indexable, then you change the value of that field. Performance is slightly slower than if the field were not indexable. This is true for any object of the class, whether or not the object is in a collection.

Now suppose an object with an indexable field is in a collection and the collection has an index on the indexable field. If you change the value of the field, performance is slightly slower than when the object is not in a collection. The extra time is needed to update the index.

An object can belong to many collections. Each collection can have an index on a particular field. If you change the value of that field, `ObjectStore` must update each index and the performance penalty is greater.

If a class has any indexable fields, every instance of the class is larger by three 32-bit words in the database.

```
-inplace
```

Causes the postprocessor to annotate stand-alone files — files that are not in `.zip` files or `.jar` files — in place rather than writing the annotated file in the destination directory. When the postprocessor annotates a class in place, it overwrites the original class files with the annotated class files and writes the `ClassInfo` subclass to the same directory as the persistence-capable class. If a class originates in a `.zip` file or `.jar` file, the postprocessor writes the annotated class and its corresponding `ClassInfo` subclass to the destination directory.

Do not use this option when you are doing iterative development. During development, a separate output directory avoids errors and supports debugging.

When you specify the `-inplace` option, you must still specify a destination directory, but the postprocessor ignores it for stand-alone files.

```
{ -it | -ignoretransient } field_name
```

Instructs the postprocessor to ignore the transient attribute of the specified field and treat the field as a persistence-capable field. You must specify a fully qualified field name. The field is treated as persistence capable only for purposes of postprocessing. This option is useful when a persistence-capable class you are defining inherits from a class that includes a transient field. If you do not specify this option for a transient field, the postprocessor ignores the field, which can cause problems if you want to use the field.

```
-modifyjava
```

Allows the postprocessor to modify classes in standard Java packages. The default is that the postprocessor does not modify standard Java classes.

```
{ -naf | -noannotatefield } qualified_field_name
```

Prevents access to the specified field from causing `fetch()` and `dirty()` calls on the containing object. This is useful for transient fields when you access them outside a transaction. Normally, access to a transient field causes `fetch()` or `dirty()` to be called to allow the `postInitializeContents()` and `preFlushContents()` methods to convert between persistent and transient state.

```
{ -natf | -noannotatetransientfields }
```

Prevents the postprocessor from inserting `fetch()` or `dirty()` calls into the annotated code for transient fields of persistent classes that are

accessed by methods. This option is useful when you want to access the transient fields outside of transactions.

The default behavior, if the `-noannotatetransientfields` option is not specified, is for the postprocessor to insert `fetch()` or `dirty()` calls into the annotated code when transient fields of persistent classes are accessed.

`-noarrayopt`

Disables optimization of `fetch()` and `dirty()` calls for array objects in looping constructs. This causes `osjcfp` to insert the calls to `fetch()` or `dirty()` in every iteration rather than only in the first loop iteration.

`{-nodefaulthashCode | -ndhc }`

Prevents the postprocessor from automatically adding a `hashCode()` method to a class, except classes for which you explicitly specify the `-hashCode` option. If you specify this option, it is your responsibility to ensure that there is a suitable `hashCode()` method for classes that are used as keys in persistent hash tables.

`-noinitializeropt`

Disables optimization of `fetch()` and `dirty()` calls in constructors. Specify this option when you want the postprocessor to perform full annotation of constructors. Full annotation means that if the object becomes persistent during constructor execution, modifications to the object are handled correctly. By default, the postprocessor does not fully annotate constructors to handle changes in the newly constructed object. Typically, this is the desired behavior.

If your application inserts objects into `ObjectStore` collections during construction of the objects being inserted, you must specify the `-noinitializeropt` option. Doing so avoids errors in the handling of modifications to the newly constructed objects.

`-noopt`

Disables the three optimizations that are disabled by the `-noarrayopt`, `-noinitializeropt`, and `-nothisopt` options. The `-noopt` option is a shortcut you can use when you want to specify all three options. You might want to specify this option when the optimizations are preventing the postprocessor from inserting required `fetch()` and `dirty()` calls in your classes.

`-nooptimizeclassinfo`

Instructs the postprocessor to generate `xxxClassInfo` classes for persistence-capable classes that are public. Use this option if your application encounters run-time security errors when the `xxxClassInfo` classes are dynamically generated using the reflection API.

`-nothisopt`

Disables optimization of `fetch()` and `dirty()` calls for access to fields relative to `this` in nonstatic member methods. This causes `osjcfp` to insert a `fetch()` or `dirty()` call for each access to a field in `this`.

`-nowrite`

Performs process and error checking but does not actually annotate class files. This option allows a test run of the postprocessor. You use it to determine whether or not all specified classes are accessible, whether additional options are needed, and if you specify `-v` (verbose), you can see where the resulting files would be located.

{ `-pa` | `-persistaware` }

Causes subsequent `.class` files, `.jar` files, and `.zip` files on the command line to be persistence aware. This means that instances of the classes can operate on persistent objects but cannot be persistent. The postprocessor annotates persistence-aware classes so that there are calls to `ObjectStore.fetch()` and `ObjectStore.dirty()` where needed during operations on potentially persistent objects and arrays that might be used by the persistence-aware class. This option applies to class names, `.class` files, `.jar` files, and `.zip` files that you specify on the command line after the `-persistaware` option and before the next `-persistcapable` or `-copyclass` option, or the end of the command line. In other words, the `-persistcapable` option or the `-copyclass` option alters this mode. The `-pc` option is in effect by default.

{ `-pc` | `-persistcapable` }

Causes subsequent `.class` files, `.jar` files, and `.zip` files on the command line to be persistence capable. This option applies to class names, `.class` files, `.jar` files, and `.zip` files that you specify on the command line after the `-persistcapable` option and before the next `-persistaware` or `-copyclass` option or the end of the command line. The `-pa` (persistence-aware) option or the `-copyclass` option alters this mode. The `-pc` option is in effect by default.

{ `-q` | `-quiet` }

Causes the postprocessor to refrain from displaying warnings. A warning message provides information about something that the postprocessor recognizes as a possible problem but cannot confirm as actually being a problem. This option cancels a previous `-verbose` option, if you specified one.

```
{ -qc | -quietclass } class_name
```

Causes the postprocessor to refrain from displaying warnings for the specified class. A warning message provides information about something that the postprocessor recognizes as a possible problem but cannot confirm as actually being a problem. This option applies only to the name that immediately follows it. Specify a fully qualified class name. If the postprocessor is renaming the class, it does not matter whether you specify the old name or the new name. If you specify `-verbose` in the same command, this option takes precedence over the specified class.

```
{ -qf | -quietfield } member_name
```

Causes the postprocessor to refrain from displaying warnings for the specified class field. A warning message provides information about something that the postprocessor recognizes as a possible problem but cannot confirm as actually being a problem. This option applies only to the name that immediately follows it. Specify a fully qualified class field name. If the postprocessor is renaming the class, it does not matter whether you specify the old name or the new name. If you specify `-verbose` in the same command, this option takes precedence for the specified class field.

```
-summary gen_class_name
```

Causes the postprocessor to generate a class with the name *gen_class_name*. This generated class extends the `com.odi.PersistentTypeSummary` class. The *gen_class_name*.`getPersistentClasses()` method returns a list of the classes that were made persistence capable in this execution of the postprocessor. The generated class has a no-argument constructor that passes arrays to the `PersistentTypeSummary` class constructor. To identify persistence-capable classes in libraries and include them in this summary, specify the `-includesummary` option.

```
{ -tf | -transientfield } qualified_field_name
```

Causes the postprocessor to treat the specified field as though it has a `transient` modifier, even if it does not. This typically is useful when a

field should not be stored in a database, but it must be available for object serialization.

```
{ -tp | -translatepackage } orig_pkg_name new_pkg_name
```

Renames classes that belong to *orig_pkg_name* so that they belong to *new_pkg_name*. The original *.class* files remain in the original location and the postprocessor does not annotate them. For example, suppose the postprocessor makes a class named *a.b.C* persistent with `-tp a.b a.b.x`. The persistent class has the name *a.b.x.C*.

A package specification of `"."` implies the default unnamed package. For example, the option `-tp . persist` causes the unpackaged class name *C* to be renamed *persist.C*.

orig_pkg_name must exactly match the package name of the class being annotated. For example, for a file named *a.c.D*, a specification of `-tp a a.b` does not translate the package name. The package of *a.c.D* is *a.c*, not *a*.

The postprocessor changes the package name of all classes in the original package that it can locate through the `CLASSPATH` environment variable or, if it is specified, the `-classpath` option.

```
{ -v | -verbose }
```

Causes the postprocessor to write descriptions of its actions to standard output. This option cancels a previous `-quiet` option, if you specified one.

File Names

You can specify any number of class names, *.class* files, *.jar* files, or *.zip* files on the command line. The postprocessor recognizes files as follows:

```
name.class
```

Explicit file name for a class file.

```
name.zip
```

Explicit file name for a class *.zip* file. The postprocessor processes all *.class* files in the *.zip* file according to the persistence mode that is in effect when the postprocessor encounters the name of the *.zip* file. The postprocessor places each file from the *.zip* file in the package-relative subdirectory of the destination directory. A *.zip* file allows the postprocessor to process multiple files without the specification of each

one on the command line. Also, you can simply specify a `.zip` file. You need not unzip the file before processing the `.class` files.

`name.jar`

Explicit file name for a class `.jar` file. The postprocessor treats the file the same way that it treats a `.zip` file.

`name`

Qualified class name delimited by `"."`. The postprocessor uses the `CLASSPATH` environment variable or the specification for the `-classpath` option to locate the `.class` file, which can be in a `.zip` file or a `.jar` file.

Because the postprocessor recognizes `name.class` as well as `name`, you can run a command such as

```
osjcfp -dest osjcfpout *.class
```

You need not derive qualified class names from the file paths.

Postprocessor API

ObjectStore also includes a companion API for its class file postprocessor utility for those times when it is useful to call the postprocessor from your Java application. The method that you call is

```
com.odi.filter.OSCFP.filter(String[] argv)
```

The arguments you pass with the `argv` parameter are the same as the arguments you would use for the `osjcfp` command-line utility. These arguments are listed in the preceding sections. The following example shows how to use the postprocessor API:

```
String args = new String[numArgs];
args[0] = "-inplace"
args[1] = "-dest";
args[2] = ".";
...;
filter.OSCFP myPostprocessor = new filter.OSCFP();
myPostprocessor.filter(args);
```

Note

When using this method, you cannot use wildcards with file names. For example, on the command line, `*.class` expands to all the `.class` files listed in your current directory; however, it does *not* expand when used with the API. You must specify each file name explicitly.

osjcggen: Generating Peer Classes

To map C++ classes to Java peer classes, run the peer generator tool (osjcggen).

Additional information about the peer class generator is in the book *Developing Java Applications That Access C++*.

Command-Line Format

The command-line format follows. It shows the different components in the command line on different lines only for clarity.

```
osjcggen -package destination_package
        -native_interface interface_type
        -schema mapping_schema_database
        -classdir target_dir -libdir target_dir [options] class_list
```

Components

`-package destination_package`

Identifies the package in which the peer generator creates peer classes. Required.

`-native_interface interface_type`

Specifies the virtual machine interface for which the tool should generate C++ glue code. Specify `jni` for the Sun VM or `ms_raw` for the Microsoft VM.

`-schema mapping_schema_database`

Identifies the mapping schema database from which the tool extracts type definitions. Required.

You must have specified the `-store_member_functions (-smf)` and `-store_function_parameters (-sfp)` options when you ran `ossd` to create this schema. If you did not, the tool cannot define peer methods. These options cause additional information to be stored in the schema database. `osjcggen` needs this information to process member functions correctly.

The mapping schema database you specify is not the application schema database that the application uses.

If the C++ code for which you want to define peer classes does not define methods, you need not specify the `-smf` and `-sfp` options when you run `ossq`.

`-classdir target_dir`

Specifies the directory relative to which the tool writes the peer classes and other generated Java files. For example, if the specified `destination_package` is `MyPkg` and the specified `target_dir` is `MyDir`, the tool creates the peer classes in `MyDir/MyPkg`. The directory you specify must exist. If the directory you specify does not contain a subdirectory with the name of the package, the tool creates this subdirectory. Required.

`-libdir target_dir`

Specifies the directory in which the tool creates the C++ files. In the directory you specify, the tool creates a directory that has the name that you specify for `destination_package`. The tool places the C++ files in this package subdirectory of the directory you specify with the `-libdir` option. The directory you specify must exist. Required.

options

Any of the options described in the next section. Optional.

class_list

Lists the names of the classes and structs for which you want the tool to generate peer classes. The tool places each peer class in the specified destination package. You can intersperse class name specifications with the options described in the next table. Required.

Do not specify `enums`. If a peer method accepts or returns an `enum`, the tool automatically generates the peer class for the `enum`.

Additional Options

The following describes the additional options you can specify on the `osjcggen` command line. These options can precede, follow, or be interspersed with the required `osjcggen` options.

`-boolean typedef_name`

Indicates that the specified `typedef` name identifies a `boolean` data type. You can specify this option multiple times.

`-classpath class_path`

Instructs the tool to use the `class_path` you specify rather than the setting for the `CLASSPATH` environment variable. When you specify this option, it

affects only the lookup of classes by the peer generator tool. It has no effect on the execution of the peer generator tool itself.

`-full`

Turns off the `-leaf` option. The tool generates complete peer classes for types specified subsequently. The `-full` option is in effect by default.

`-import package`

Specifies a package name in which to look for a peer class before creating it. For example, suppose you want to generate a peer class for the C++ class `x`. The tool looks in the package specified with the `-import` option for a Java version of `x`. If such a version is found, the tool does not generate a new peer class for `x`. The application can use the existing peer class. You can specify this option multiple times.

`@input_file`

You can use the `@input_file` option to specify a file that contains arguments for the peer generator tool. The tool inserts the contents of the specified file in the command line before it begins to execute the command line. You can specify this option multiple times. An input file cannot itself include the `@input_file` option. If it does, the tool treats it as the name of a class, which is not found.

`-leaf`

Instructs the tool to generate a minimal definition for the peer classes of types specified after this option and before any `-full` option. When the tool generates a leaf peer class, the application cannot access C++ data or function members or any C++ base classes. The tool defines the Java peer class to inherit directly from `CPlusPlus`. It does not copy the inheritance structure from the C++ class; this option is useful when you want to prune a type graph to reduce the size of the interface code.

`-map C++_name Java_peer_name`

Allows you to specify the name of the Java peer class that you want to identify a particular C++ class. When you specify this option, you need not rely on the default name mapping rules. This option is particularly useful for naming template classes. You can specify the `-map` option as many times as you need to. Follow each specification with

- a The name of a C++ class. You must also specify the name of the C++ class in the list of classes for which you want to generate peer classes.

- b** The name of the Java peer class that you want to represent the C++ class. This can be a fully qualified class name or an unqualified class name. If it is unqualified, the peer generator tool places the Java peer class in the package you specify on the command line.

`-map_existing C++_name Java_peer_name`

Allows you to specify the name of an existing Java peer class that you want to identify a particular C++ class. This option is the same as the `-map` option, except that the Java peer class already exists or will exist. Consequently, the peer generator tool does not generate any peer code for the Java peer type. In a separate run of the peer generator tool, or manually, you must create the Java peer type. See the documentation for `-map` for additional details.

`-nosynchronize`

Turns off automatic synchronization of peer methods. If you specify this option, your application is responsible for ensuring that only one thread at a time per session is accessing ObjectStore. Failure to prevent concurrent access to OSJI and peer method entry points can cause OSJI to fail.

`-oldtemplates`

Causes the peer generator to map some C++ characters to Java equivalents used in previous releases of ObjectStore. Rules for Template Name Flattening in Chapter 2 of *Developing Java Applications That Access C++* shows the C++ characters that are invalid in Java and the Java equivalents that they are mapped to.

`-synchronize`

Turns on automatic synchronization of peer methods. This is the default.

`-suppress package.class.method`

Suppresses generation of the specified peer method. You should not need to specify this option frequently. However, if generated code for a particular method causes a problem for the compiler, this option allows you to prevent generation of that code.

You must specify the package name, the class name, and the method name. Use this option in the following manner:

- c** Run `osjcggen` without specifying this option.
- d** Try to compile the generated code.
- e** A particular peer method causes a problem.

- f Run `osjcggen` again and specify the same classes, but suppress generation of the peer method that causes the problem.
- g Determine how to work around the lack of the suppressed peer method.

If it is not possible to work around the method that is causing the problem, you must redefine the C++ method into a form that does not cause a problem when it is mapped to a peer method.

Example of Running the Peer Generator Tool

The following example generates a Java peer class for the `CPerson` class. If it does not already exist, the tool creates the `MyPkg` subdirectory in the `MyDir` directory. The tool puts the Java peer class in the `MyPkg` package/subdirectory. The mapping schema database that the tool uses to generate the peer class is `CPerson.jadb`. The tool places the generated C++ files and the generated Java files in the `MyDir/MyPkg` directory.

```
osjcggen -package MyPkg
        -native_interface jni
        -schema CPerson.jadb
        -classdir MyDir
        -libdir MyDir
CPerson
```

If you specify

```
-libdir SomeOtherDir
```

the tool places the generated Java files in `MyDir/MyPkg` and the generated C++ files in `SomeOtherDir/MyPkg`.

Caution

If the tool generates files that have the same pathnames as your C++ source files, the tool overwrites the C++ source files without warning you. This can happen if you create your C++ source files in the package and directory that you specify with the `-package`, `-classdir`, or `-libdir` options to `osjcggen`.

For example, suppose you have the `CPerson.cc` C++ file in the `/MyDir/MyPkg` directory. When you run `osjcggen`, you specify the following on one line:

```
osjcggen
        -package MyPkg
        -native_inteface jni
        -schema CPerson.jadb
        -classdir MyDir
```

```
-libdir MyDir  
CPerson
```

The peer generator tool generates the following files:

```
/MyDir/MyPkg/CPerson.java  
/MyDir/MyPkg/CPersonU.java  
/MyDir/MyPkg/CPersonClassInfo.java  
/MyDir/MyPkg/CPerson.cc  
/MyDir/MyPkg/CPersonU.cc
```

An explanation of these files appears in Output from the Peer Generator Tool in *Developing Java Applications That Access C++*. But as you can see, the tool would overwrite your `CPerson.cc` source file. On Solaris, you can work around this by specifying `.CC` instead of `.cc` in the C++ file name. On Windows, you can work around this by specifying `.cpp` instead of `.cc` in the C++ file name. Alternatively, you can either specify different directories for `-classdir` or `-libdir`, or you must not create C++ source files in the Java package subdirectory.

osjcheckdb: Checking References in a Database

The `osjcheckdb` utility or the `Database.check()` method checks the references in a database. This tool scans a database and checks that there are no references to destroyed objects. You can fix references to destroyed objects by finding the objects that contain the dangling references and overwriting the invalid references with something else, such as a null value. In addition to finding references to destroyed objects, the tool performs various consistency checks on the database.

If the tool does not find any problems, it does not produce any output.

Check paths Before you invoke `osjcheckdb` from the command line, ensure that `tools.jar` is in your `CLASSPATH` variable. Also ensure that the distribution `bin` directory that contains `osjcheckdb` is in your `PATH` variable. The format for invoking this tool from the command line is

Command line `osjcheckdb database_name1.odb ...`

You can specify one or more databases. Separate multiple specifications with a space. If `osjcheckdb` cannot check a database that you specify, it displays a message about the inaccessible database and continues to the next database.

The tool displays messages about any errors that it finds.

API

The function signature for invoking the API for this tool is

```
Database.check(java.io.PrintStream)
```

When `ObjectStore` executes this method, it operates on the committed contents of the database and on any changes that have been saved as a result of `ObjectStore.evict()` or `ObjectStore.evictAll()`. `ObjectStore` does not operate on any changes that have been made but not committed or evicted.

The method writes any errors it finds to the argument stream. It also returns a Boolean value, which is `true` if the references are valid and `false` if there are any bad references. The `osjgcdb` tool requires

- The `tools.jar` file in your `CLASSPATH` environment variable
- The `ObjectStore bin` directory in your `PATH` environment variable

osjshowdb: Displaying Information About a Database

The `osjshowdb` utility displays information about one or more databases. This utility is useful when you want to know how many and what types of objects are in a database. You can use this utility to verify the general contents of the database.

This utility displays the following information:

- Name of the database
- Size of each segment
- Name and number of each type of object in the database
- Total size in bytes occupied on the disk by each type of object

Number of
destroyed
objects -
`showObjs`
option

If you specify the `-showObjs` option, the `osjshowdb` utility also displays the following information for each object:

- `oid`, which is an internal representation of its location in the database
- Type
- Number of bytes it occupies on the disk
- If it is an array, the number of elements in the array

-showData option	Specify the <code>-showData</code> option with <code>osjshowdb</code> to display string values as well as the references that an object contains. When you specify the <code>-showData</code> option, it implies the <code>-showObjs</code> option.
Path variables	Before you invoke <code>osjshowdb</code> from the command line, ensure that <code>tools.jar</code> is in your <code>CLASSPATH</code> environment variable. Also ensure that the distribution <code>bin</code> directory that contains <code>osjshowdb</code> is in your <code>PATH</code> variable.
Command line	<p>To execute the <code>osjshowdb</code> utility, use this format:</p> <pre>osjshowdb [-showData] [-showObjs] db1.odb [db2.odb]...</pre> <p>You can specify one or more databases.</p> <p>When the utility displays <code>java.lang.String</code> objects, the number of elements is the number of characters in the string. The total bytes indicates the number of bytes that the data occupies on the disk.</p> <p>There are internal structures in the database that are not included in the calculations performed by the <code>osjshowdb</code> utility. Consequently, the total number of bytes as indicated in the output from <code>osjshowdb</code> is never equal to the actual size of a segment.</p>
API	<p>The API for the <code>osjshowdb</code> utility is <code>Database.show()</code>.</p> <p>When <code>ObjectStore</code> executes this method, it operates on the committed contents of the database and on any changes that have been saved as a result of <code>ObjectStore.evict()</code> or <code>ObjectStore.evictAll()</code>. <code>ObjectStore</code> does not operate on any changes that have been made but not committed or evicted.</p>

osjuphsh: Upgrading String Hash Codes in Databases

The `osjuphsh` utility is the command line utility for upgrading databases to use the correct `String` hash code. Invoke this tool with the following format:

```
osjuphsh [options] database_name
```

For example:

```
osjuphsh db.odb
```

Execution of this command is the same operation as a call to the `com.odi.Upgrade.upgradeDatabaseStringHash()` method on the `db.odb`

database with the `upgradeObjects` argument set to true and the `multipleTransactions` argument set to false.

The options you can specify are listed in the following table:

<i>Option</i>	<i>Description</i>
<code>-skipobjects</code>	Specifies that the database contains no objects that depend on <code>String.hashCode()</code> .
<code>-multitrans</code>	Specifies that each object should be upgraded in a separate transaction.
<code>-verbose</code>	Specifies that verbose information about the upgrade should be displayed. This includes displaying information about databases that do not require upgrading.

osjversion: Obtaining ObjectStore Version Information

You can use the `osjversion` utility to display the version number and the build date for the version of ObjectStore you are running. This command is in the `bin` subdirectory of the installation directory. You must include `bin` in your path to run `osjversion`. For example:

```
% osjversion
ObjectStore Java Interface Release 6.0 (1999-03-15 12:00:00
on rusa)
ObjectStore Release 6.0 for Windows NT and Windows 98 Systems
Build packaged 03-15-99 12:00
%
```

This command is useful when you want to ensure that you have the right version in your path.

You can also run the `unzip -z` command against any ObjectStore `.jar` file. Doing so allows you to see the version of ObjectStore that the `.jar` file contains. For example:

```
% unzip -z osji.jar
Archive: osji.jar ObjectStore Java interface 6.0 (1999-03-15
12:00:00 on rusa), Copyright 1996-1999, eXcelon Corporation,
```

Inc. All rights reserved.

%

The `osjversion` utility checks what is available through the `PATH` environment variable. To check what is available through the `CLASSPATH` variable, you can write a program such as the following.

```
import com.odi.ObjectStore;
class OSVersion {
    public static void main (String[] args) {
        ObjectStore.initialize(null,null);
        System.out.println(ObjectStore.releaseName());
    }
}
```

Appendix

Packaging Your Application for End Users

When you package your application for delivery to end users, the package must include one and possibly two class files for each persistence-capable class in your application:

- The annotated class file
- The corresponding `ClassInfo` class file, if one was generated

For example, if you have a persistence-capable class called `Person`, in the `App` package, you must provide users with

- The annotated `App\Person.class` file
- The `App\PersonClassInfo.class` file, if present

There is no corresponding `ClassInfo` file for persistence-capable interfaces.

The class file postprocessor only generates `ClassInfo` files for persistence-capable classes if the class does not provide public access to all the required fields and methods. In most cases, `ClassInfo` instance is created as needed at run time. Check the destination directory specified for the postprocessor to determine which `ClassInfo` files were generated.

You can jar these files with the rest of your package. You need not send the unannotated version of your persistence-capable classes. The annotated version can be used in a transient context.

For persistence-aware classes, you must provide the annotated class file. There is no corresponding `ClassInfo` class file.

Library providers

If you are providing a library that contains persistence-capable classes, you must also provide the `PersistentTypeSummary` object that identifies those classes. `ObjectStore` uses the summary object to install schema information for your library at run time. This allows users to install a new version of a library without rebuilding the application type summary.

Glossary

active persistent object

An active persistent object starts as an exact copy of the object that it represents in the database. ObjectStore initializes a hollow object so that it becomes an active object. This happens when an application calls the `ObjectStore.fetch()` or `ObjectStore.dirty()` method. If an application calls the `ObjectStore.dirty()` method, rather than the `ObjectStore.fetch()` method, on a hollow object, the resulting active object can be modified.

Consequently, an active object is not necessarily identical to the object in the database that it represents. An application can read or update an active persistent object; a persistent object must be active for an application to read or update it.

affiliated database

A database whose objects can be referenced by objects in another database.

API object

An object defined by the ObjectStore API that is associated with one session. An API object can be an object of one of the following classes: `Cluster`, `Database`, `DatabaseRootEnumeration`, `DatabaseSegmentEnumeration`, `Segment`, or `Transaction`.

batch

A batch is a set of files that must be postprocessed together. Often, this is all the files in your application. In more complex applications, there might be multiple batches that each contain a library and a batch of files that you write, which reference the libraries.

database

Persistent storage is organized into databases. Before a persistent object can be created, the database in which it is to be stored must exist, and this database must be opened by the process performing the creation. The database must also be opened by any processes accessing the object. A single application can open several databases simultaneously. A single database can be accessed by many applications simultaneously.

deadlock	A simple deadlock occurs when one transaction holds a lock on a data item that another transaction is waiting to access, while the second transaction holds a lock on a data item that the first transaction is waiting to access. Neither process can proceed until the other does. <code>ObjectStore</code> detects and breaks deadlocks by aborting one of the transactions involved.
hollow persistent object	A hollow persistent object contains fields that are identical to the fields of the object in the database that the persistent object represents, but the fields have default values.
multistep index	An index on a complex navigational path that accesses multiple public data members, such as <code>a.b().c.name</code> .
persistence-aware	<p>If the methods of a class can operate on persistent objects but an instance of the class cannot itself be stored in a database, the class is persistence aware.</p> <p>When a method accesses fields in a persistent object, <code>ObjectStore</code> checks to ensure that the data has been read from the database. This checking is done by calls to the <code>ObjectStore.fetch()</code> and <code>ObjectStore.dirty()</code> methods. A persistence-aware class includes the annotations that call the <code>fetch()</code> and <code>dirty()</code> methods. It does not include the other annotations required for a class to be persistence capable. Normally, you run the class file postprocessor to annotate a class so that it is persistence aware. Occasionally, you manually annotate the class yourself.</p>
persistence-capable	<p>A persistence-capable object has the capacity to be stored in a database. If you can store the instances of a class in a database, the class is persistence capable and the instances of the class are persistence-capable objects.</p> <p>The definition of a persistence-capable class includes specific annotations required by <code>ObjectStore</code>. After you compile class definitions, you run the <code>ObjectStore</code> class file postprocessor on the compiled classes to add the annotations that make the classes persistence capable.</p> <p>Some Java-supplied classes are persistence capable. Others are not persistence capable and cannot be made persistence capable. A third category of classes can be made persistence capable but there are important issues to consider when you do so. Be sure to read <i>Java-Supplied Persistence-Capable Classes</i> on page 311.</p>
persistent object	<p>A persistent object is a representation of an object that is stored in a database.</p> <p>After an application retrieves an object from the database, the application works with the persistent object in the Java environment. A persistent object always exists in one of three states:</p>

- Hollow
- Active
- Stale

session

A session allows the use of the ObjectStore API. ObjectStore uses the abstract `com.odi.Session` class to represent sessions.

Your application must create a session. After a session is created, it is an active session. A session remains active until your application or ObjectStore terminates it. After a session is terminated, it is never used again. You can, however, create a new session.

A session consists of a set of persistent objects and a set of ObjectStore API objects such as a `Transaction`, `Databases`, and `Segments`.

In a single Java VM, PSE Pro and ObjectStore allow multiple concurrent sessions.

If you are using ObjectStore or PSE Pro, separate Java virtual machines can each run multiple sessions at the same time. See *How Sessions Keep Threads Organized* on page 19.

**stale
persistent
object**

A stale persistent object is no longer valid. Its fields have default values and it should not be used.

A persistent object might become stale after an application commits or aborts a transaction in which the active or hollow persistent object was accessible. When an application calls the `ObjectStore.destroy()` method, the target of the destroy method becomes stale.

An application must not try to read or update a stale object.

**transitive
persistence**

When an application commits a transaction, it stores in the database any transient objects that can be transitively reached from any persistent objects. This is the process of transitive persistence.

Index

A

- `AbortException` 93
- `AbortException.getOriginalException()` method 122
- aborting transactions 92
 - default effects on persistent objects 146
 - setting default object state 139
 - setting objects to default state 140
 - specifying a particular object state 140
- abstract classes 252
- accessing persistent objects
 - committing transactions 122
 - default effects of methods 146
 - `dirty()` method 242
 - evicting objects 132
 - `fetch()` method 242
 - optimizing 230
 - procedure 102
 - saving changes by committing transaction 122
 - saving changes through eviction 132
- active persistent objects
 - aborting transactions 141
 - committing transactions 127, 130
 - default effects of methods 146
 - definition 8
 - evicting 135
- adding thread to session 30
- `addressSpaceSize` property 38
- `addSuperIndex()` method in JDD 306
- aggregations, very large 161
- annotations
 - customizing 226
 - description 197
 - manual 237
 - superclass modifications 214
 - you must add 234
- `applicationName` property 39
- applications
 - complex, finding right class files 210
 - failure 3
 - required components 13
- architecture 4
- archive logging 2
- arrays
 - optimizations 328
 - passing 235
- attributes, JDD
 - assigning values to 296
 - overview 292
 - type safety 292

B

- Bad command or file name error message 217, 322
- batch

- definition 198
- postprocessing two 198
- postprocessor requirement 201

batch schema installation

- advantage 272
- database, creating with 275
- identifying application types 273
- installing application types 276
- introduction 271
- procedure 273

-boolean option 334

breakpoint

- debugging 218

BrokenServerException 93

C

C++ applications 2

C++ database garbage collection 64

cache manager

- description 4

cache size 40

cachedObjectCount property 39

cachedObjectTransactionAge property 39

cacheSize property 40

changing classes

- schema evolution 65

checkpoint 264

checkpointing transactions

- setting default object state 266

Class could not be found error 202

class file

- transient version 203

class files

- annotated
 - finding 210
 - locations 215
 - managing 208
- applications, complex 210
- compiling unannotated 209
- inner 207
- nested 207
- referring to persistent and transient versions 223

ClassCastException troubleshooting 46

-classdir option 334

ClassInfo class

- classinfosuffix option 324
- nooptimizeclassinfo option 207
- subclass, defining 243
- when generated 207

-classinfosuffix option 324

ClassNotRegisteredException 196

CLASSPATH environment variable

- alternatives 211
- class files, locating 209
- classpath option 204
- requirements 200

-annotatefield option 324

-classpath option 204, 324, 334

clearContents() method

- postprocessor 226

client description 4

Cluster.acquireLock() method 268

Cluster.getObjects() method 108

clustering 259

clusters

- creating 54
- description 54
- grouping objects 78
- objects, iterating through 55, 108
- planning 78
- storing objects in 55
- storing objects in particular ones 101

collections

- adding indexes 185
- advantages 166
- alternative, selecting best 168
- bags 160
- built-in types, storing 193
- choosing 168
- comparison 168

- creating 171
- hash code 159
- hash code requirements 192
- implemented interfaces 159
- inserting objects during
 - construction 219, 328
- introduction 157
- iterating 171
- lists 165
- navigating 171
- OSHashBag 160
- OSHashMap 160
- OSHashSet 160
- OSHashtable 161
- OSTreeMap 161
- OSTreeSet 163
- OSVector 165
- OSVectorList 165
- postprocessor optimization 219
- querying 173
- relative size 168
- sets 160
- third-party 194
- com.odi.addressSpaceSize property 38
- com.odi.applicationName property 39
- com.odi.cachedObjectCount property 39
- com.odi.cachedObjectTransactionAge
 - property 39
- com.odi.cacheSize property 40
- com.odi.disableObjectCaching
 - property 40
- com.odi.disableWeakReferences
 - property 41
- com.odi.gc.displayGarbage property 64
- com.odi.gc.lockTimeOut property 63
- com.odi.gc.reachableObjects
 - property 64
- com.odi.gc.reclaimedObjects
 - property 64
- com.odi.gc.retries property 63
- com.odi.gc.retryInterval property 63
- com.odi.gc.transactionPriority
 - property 63
- com.odi.migrateUnexportedStrings
 - property 42
- com.odi.ObjectStore.library
 - property 42
- com.odi.password property 42
- com.odi.product property 42
- com.odi.Session class 20
- com.odi.stringPoolSize property 45
- com.odi.trapUnregisteredType
 - property 46
- com.odi.useImmediateStrings
 - property 46
- com.odi.user property 42
- com.odi.util.query.Query class 173
- commit(ObjectStore.RETAIN_
 - READONLY) 127, 130
- commit(ObjectStore.RETAIN_UPDATE)
 - method 129
- committing transactions
 - after evicting objects 137
 - default effects on persistent objects 146
 - failures 91
 - introduction 90
 - RETAIN_HOLLOW 126
 - RETAIN_READONLY 127
 - RETAIN_STALE 125
 - RETAIN_TRANSIENT 130
 - RETAIN_UPDATE 129
 - saving changes 122
 - setting default object state 124
 - setting objects to default state 125, 267
 - situations to avoid 101
- compacting a database 64
- concurrency control
 - access, preventing 267
 - batch schema installation 271
 - conflicts, handling 270
 - locking objects 267
 - multiversion 261

- consistent state 96
- constructors
 - postprocessor optimization 219
- cooperating threads 31
- copyclass option 204, 324
- copying classes without annotating 204
- copying data between PSE Pro and
 - ObjectStore databases 42
- copying databases 61
- creating clusters 54
- creating database roots 104
- creating databases 50
- creating external references 110
- creating notifications 282
- creating sessions 23

D

- Database class
 - description 50
- database operations
 - table of 81
- database roots
 - changing object referred to 106
 - creating 104
 - destroying 107
 - how many 107
 - null values 106
 - primitive values 106
 - retrieving 105
- Database.acquireLock() method 269
- Database.check() method 338
- Database.close() method 58
- Database.create() method
 - batch schema installation 275
 - default schema installation 50
 - example 50
- Database.createRoot() method 104
- Database.createSegment() method 53
- Database.destroy() method 75
- Database.destroyRoot() method 107
- Database.evolveSchema() method 67
- Database.GC() method 62
- Database.getAffiliatedDatabases()
 - method 77
- Database.getOpenMode() method 76
- Database.getPath() method 76
- Database.getSegments() method 54
- Database.getSizeInBytes() method 76
- Database.installTypes() method 276
- Database.isOpen() method 75
- Database.open() method 57
- Database.setRoot() method 106
- databases
 - autoopen 60
 - closing 56
 - clusters 54
 - compacting 64
 - consistent state 96
 - copying 61
 - copying data between PSE Pro and
 - ObjectStore 42
 - creating 50
 - creating roots 104
 - destroying 75
 - destroying objects 142
 - displaying information about 339
 - dumping and loading 73
 - exporting objects 80
 - garbage collection 62
 - identity 59
 - information about, displaying 77
 - information about, obtaining 75
 - locking with acquireLock() 267
 - managing 49
 - moving 61
 - MVCC, open for 262
 - names 52
 - objects, migrating 78
 - objects, storing 100
 - open modes 60
 - open types 57
 - open? 75

- opening 56
 - opens, automatic 60
 - pathname of 76
 - platforms 50
 - read-only, open for? 76
 - references, checking 338
 - roots, how many 107
 - schema evolution 65
 - segments 53
 - size of 76
 - transient 56
 - update, open for? 76
 - upgrading 73
 - dead objects 339
 - DeadlockException 93
 - deadlocks
 - aborting transactions 96
 - description 92
 - retrying aborted transactions 94
 - debugger 217
 - deepFetch() method
 - description 129
 - serialization 317
 - dest option 325
 - destination directory
 - about 202
 - requirement 201
 - destroying database roots 107
 - destroying databases 75
 - destroying objects
 - cleaning up references 145
 - ObjectNotFoundExceptions 145
 - persistent objects, default effects 146
 - preDestroyPersistent() hook
 - method 142
 - destroying objects in the database 142
 - destroying objects referred to by other objects 145
 - dirty() method
 - background 253
 - manual annotations 242
 - disableObjectCaching property 40
 - disableWeakReferences property 41
 - disk space
 - copy of object in database 119
 - dumping databases 73
 - duplicates
 - postprocessor, file specifications for 205
 - strings 312
 - dynamic data, classes for modelling 291
- ## E
- embedded strings 325
 - embeddedmaxlengthfield option 325
 - entities, JDD
 - adding to type extent 296
 - creating 295
 - overview 293
 - enumerating segments 54
 - enums
 - specification 334
 - environment variables
 - OS_JAVA_VM 322
 - OSJCFPJAV 322
 - evicting objects
 - all 136
 - persistent objects, default effects on 146
 - persistent objects, references to 133
 - RETAIN_HOLLOW 134
 - RETAIN_STALE 134
 - threads, cooperating 136
 - outside transactions 137
 - transactions, committing 137
 - evolveSchema() method 67
 - evolving schema
 - introduction 65
 - not required for JDD 292
 - when required 66
 - examples
 - annotations, manual 244
 - basic JDD tasks 297
 - batch schema installation 276

- before running a program 17
- code for people demo 14
- compiling 17
- general use 13
- getFields() method 255
- hook methods 227
- identity 119
- indexes on collections 187
- JDD relationships 303
- persistence mode options, multiple 205
- postprocessing batches 198
- postprocessor command line 201
- querying utility collections 174
- running a program 18
- running peer generator tool 337
- running postprocessor 202
- schema evolution 70
- serialization 70
- serializing 317
- transient-only fields 249
- exceptions
 - ObjectException 80
- Exported objects
 - ObjectStore.RETAIN_STALEargument value 80
- exporting objects 80
- exporting peer objects 80
- exporting primary objects 80
- extent, adding entities to 296
- external references
 - creating 110
 - encoding as strings 112
 - introduction 109
 - obtaining objects 112
 - reusing 115
- ExternalReference class 109
- ExternalReference.fromString()
 - method 112
- ExternalReference.getObject()
 - method 112
- ExternalReference.toString()

- method 112

F

- failover 2
- fetch() method
 - background 253
 - manual annotations 242
- Field class 254
- fields in manual annotations 256
- final fields
 - initialization 225
- final fields
 - postprocessor handling of 215
- finalize() method
 - annotations 218
 - avoiding 148
- flushContents() method
 - postprocessor 226
- force option 325
- free variables 181
- full option 335

G

- garbage collection
 - active objects from commit() 128
 - active objects from evict() 135
 - C++ databases 64
 - databases 62
 - hollow objects from commit() 126
 - hollow objects from evict() 134
 - osgc utility 64
 - persistent, overview 61
 - properties 63
 - segments 63
 - stale objects from commit() 125
 - stale objects from evict() 134
 - strings 62
 - tombstones 62
 - weak references 122
- GenericObject class

- description 253
- getting field values 255
- setting field values 255
- `get.PersistentClasses()` method 273
- `getFields()` method
 - background 254
 - example 255
- global sessions 23
- grouping objects in clusters 78

H

- hash tables
 - references to destroyed objects 145
- `-hashcode` option 325
- `hashCode()`
 - arrays 193
 - requirements 192
- hollow object constructors
 - creating 229
 - transient nonstatic fields 235
- hollow persistent objects
 - aborting transactions 141
 - default effects of methods 146
 - definition 7
 - `evict()` 134
 - transactions, committing 126
- hook methods 226

I

- identity
 - databases 59
 - Java wrapper classes 312
 - persistent objects 119
- `-ignoretransient` option 327
- `-import` option 335
- `-includesummary` option
 - description 326
 - how to use 274
- `-indexablefield` option
 - description 326

- `IndexDescriptor` class 190
- `IndexDescriptorSet` class 190
- `IndexedCollection` interface 185
- `IndexedOneToMany` class in JDD 300
- indexes
 - adding in JDD 295
 - adding to collections 185
 - background 184
 - dropping 186
 - example 187
 - introduction 184
 - managing 189
 - modifying 187
 - multistep 186
 - optimizing queries 190
 - superindexes in JDD 306
 - updating 188
- `initializeContents()` method
 - postprocessor 226
- initializing API
 - specifying properties 37
- initializing objects
 - definition 7
- initializing transient fields 225
- inner classes 207
- `-inplace` option 326
- input file for postprocessor 233
- input file option 334
- `@input_file` option 323, 335
- installing application types 276
- installing schema information 271
- interfaces
 - annotations 238
- interoperability
 - specifying enums 334
- interoperability between PSE Pro and ObjectStore 42
- `IPersistentHooks` interface 241
 - hook methods 241
- iterating over entities in extent of type 297
- iterating through objects, enumerating

- clusters 55
- iterators 171

J

- jar files 206
- Java executables 217
- Java Remote Method Interface
 - See* RMI 128
- Java-supplied classes 311
- JDD
 - mixing entities and Java objects 307
 - overview 291
- JDK 1.2
 - compatibility 166

L

- large aggregations 161
- Lea, Doug, collections library 194
 - leaf option 335
 - libdir option 334
- libraries
 - application types, identifying in 274
 - ObjectStore 42
 - postprocessing existing 198
 - providing type summaries 344
 - third-party collections 194
 - type summaries, providing 274
- LinkedOneToMany class in JDD 300
- ListIterator class 171
- loading databases 73
- locking
 - databases 267
 - objects 267
 - peer objects 270
 - segments 267
 - and transaction length 260
 - wait time, reducing 259
- locks
 - acquiring 267
 - converting read to write 37

- MVCC, obtaining 262
- releasing 270
- timeouts 260
- types 269

- LockTimeoutBlocker.getApplicationName() method 270
- LockTimeoutBlocker.getHostName() method 270
- LockTimeoutBlocker.getLockType() method 270
- LockTimeoutBlocker.getPID() method 270
- LockTimeoutException.getBlockers() method 270
- long file names
 - classinfosuffix option 324

M

- manual annotations
 - abstract classes 252
 - accessing fields 253
 - application types, identifying 277
 - ClassInfo subclass definition 243
 - creating fields 253
 - example 244
 - fields, accessing 256
 - fields, creating 256
 - fields, transient-only 248
 - methods, required 239
 - persistence-aware classes 251
 - postprocessor conventions 252
 - procedure 238
- ManyToMany class in JDD 301
- many-to-many relationships in JDD 301
- ManyToManyWithObject class in JDD 302
 - map option 335
 - map_existing option 336
- mapping schema database
 - schema option 333
- maps
 - OSHashMap 160

- OSTreeMap 161
- querying 166
- media failure 3
- migrating
 - PSE Pro to ObjectStore 42
 - unexported strings 42
- modifyjava option 327
- moving databases 61
- moving objects into a database 100
- moving objects to another segment 81
- multiversion concurrency control
 - databases, multiple 262
 - databases, opening 262
 - locks, obtaining 262
 - serializability 262
 - snapshots 261

N

- native methods
 - capability for persistence 314
 - postprocessing 234
- native_interface option 333
- nested classes 207
- nested transactions 88
- noannotatefield option 225, 327
- noarrayopt option 218, 328
- nodefaultashcode option 328
- noinitializeropt option 219, 328
- noncooperating threads 32
- nonexported objects 131
- nonglobal sessions 24
- nonpersistent methods 234
- noopt option 219, 328
- nosynchronize option 336
- nothisopt option 219, 329
- notification facility
 - background 279
 - introduction 279
 - managing 287
 - performance 288
 - process flow 280

- queue 287
- security 282
- Notification() 282
- Notification.getData() method 287
- Notification.getKind() method 287
- Notification.getLocation()
 - method 287
- Notification.getMessage() method 287
- Notification.getPendingNotification
 - s() method 288
- Notification.getQueueOverflows()
 - method 288
- Notification.getQueueSize()
 - method 288
- Notification.notifyImmediate()
 - method 285
- Notification.notifyOnCommit()
 - method 285
- Notification.receive() method 286
- Notification.subscribe() method 284
- Notification.unsubscribe()
 - method 285

- notifications
 - creating 282
 - definition 280
 - reading 287
 - retrieving 286
 - sending 285
 - subscribing 284
 - threads 281
 - transactions 281
- nowrite option 232, 329
- null values
 - queries 183

O

- object caching property 40
- object table 122
- ObjectException 80
- objects
 - destroying 142

- evicting
 - See* evicting objects
- exporting 80
- external references 109
- identity 119
- is it persistent? 103
- listing in a cluster 108
- listing in a segment 108
- obtaining from external references 112
- retrieving 102, 105
- storing 100
- updating 118
- ObjectStore
 - code example of use with PSE Pro 44
 - copying data to PSE Pro 42
 - general description 1
 - process architecture 4
 - what it does 2
- ObjectStore library property 42
- ObjectStore utility collections
 - hash code method requirements 192
- ObjectStore.acquireLock() method 268
- ObjectStore.deepFetch() method 128, 317
- ObjectStore.destroy() method 142
- ObjectStore.dirty() method 242
- ObjectStore.evict() method 37, 132
- ObjectStore.evict(RETAIN_HOLLOW) method 134
- ObjectStore.evict(RETAIN_STALE) method 134
- ObjectStore.evictAll() method 136
- ObjectStore.export() method 80
- ObjectStore.fetch() method 242
- ObjectStore.getAutoOpenMode() method 60
- ObjectStore.INSTALL_SCHEMA_BATCH 275
- ObjectStore.INSTALL_SCHEMA_INCREMENTAL 275
- ObjectStore.migrate() method 79, 283
- ObjectStore.MULTI_DB_MVCC constant 57
- ObjectStore.MVCC constant 57
- ObjectStore.READONLY constant 57, 86
- ObjectStore.RETAIN_HOLLOW
 - aborting transactions 141
 - committing transactions 126
 - evicting objects 134
- ObjectStore.RETAIN_READONLY
 - aborting transactions 141
 - committing transactions 127
 - evicting objects 135
- ObjectStore.RETAIN_STALE
 - aborting transactions 140
 - committing transactions 126
 - evicting objects 134
- ObjectStore.RETAIN_TRANSIENT
 - aborting transactions 141
 - committing transactions 130
 - evicting all persistent objects 136
- ObjectStore.RETAIN_UPDATE
 - aborting transactions 141
 - committing transactions 129
- ObjectStore.setAutoOpenMode() method 60
- ObjectStore.UPDATE
 - starting transaction 86
- ObjectStore.UPDATE constant 57
- ObjectStoreConstants.MULTI_DB_MVCC 57
- ObjectStoreConstants.MVCC 57
- ObjectStoreConstants.READONLY 57
- ObjectStoreConstants.UPDATE 57
- ODMG binding 2
- oldtemplates option 336
- OneToMany class in JDD 300
- one-to-many relationships in JDD 300
- OneToOne class in JDD 300
- one-to-one relationships in JDD 300
- on-line backup 2
- optimizations, postprocessor
 - descriptions 218

- disabling 218
- optimizing JDD queries 306
- summary option
 - description 330
- options
 - @input_file 335
 - annotatefield 324
 - boolean 334
 - classdir 334
 - classinfosuffix 324
 - classpath 204, 324, 334
 - copyclass 204, 324
 - dest 325
 - embeddedmaxlengthfield 325
 - force 325
 - full 335
 - hashcode 325
 - ignoretransient 327
 - import 335
 - includesummary 274, 326
 - indexablefield 326
 - inplace 326
 - leaf 335
 - libdir 334
 - map 335
 - map_existing 336
 - modifyjava 327
 - native_interface 333
 - noannotatefield 225, 327
 - noarrayopt 218, 328
 - nodefaulthashcode 328
 - noinitializeropt 219, 328
 - noopt 219, 328
 - nosynchronize 336
 - nothisopt 219, 329
 - nowrite 232, 329
 - oldtemplates 336
 - package 333
 - persistaware 204, 329
 - persistcapable 204, 329
 - quiet 232, 329
 - quietclass 215, 330
 - quietfield 215, 330
 - schema 333
 - showData 340
 - showObjs 339
 - summary 273
 - suppress 336
 - synchronize 336
 - sysclasspath 324
 - transientfield 225, 330
 - translatepackage 221, 331
 - verbose 232, 331
- OS_JAVA_VM environment variable 322
- oscompact utility 64
- oscopy utility 61
- osdump utility 73
- osgc utility
 - C++ databases 64
 - Java databases 64
- OSHashMap collection 160
- OSHashMap collections 160
- OSHashSet collections 160
- OSHashtable collections
 - description 161
 - JDK 1.2 compatibility 167
 - lazy allocation 161
- OSJCFPJAVA environment variable 217, 322
- osjcfputility
 - command-line syntax 323
 - overview 196
- osjcgen utility 333
- osjcheckdb utility 338
- osji.jar file 17
- osjshowdb utility 77, 339
- osjversion utility 342
- osload utility 73
- osmv utility 61
- OSTreeMap collections 161
- OSTreeSet collections 163
- OSVector collections
 - description 165

JDK 1.2 compatibility 167
 OSVectorList collections 165

P

package names
 postprocessed classes 221
 renaming 236
 -package option 333
 password property 42
 PATH requirements 201
 pattern matching 179
 optimizing 180
 special characters 179
 peer generator tool
 class list 334
 example of running 337
 options 333
 running 333
 peer objects
 definition 10
 performance
 cross-segment references 77
 Java-supplied classes 314
 lazy hash table allocation 161
 lazy vector allocation 165
 notification facility 288
 -persistaware option 204, 329
 -persistcapable option 204, 329
 persistence
 how objects become persistent 100
 Java-supplied classes 311
 manual annotations 248
 transitive 100
 persistence mode options 204
 persistence-aware classes
 creating 212
 definition 9
 manual annotations 251
 persistence-capable classes
 abstract 252
 annotations 197
 definition 6
 generating automatically 195
 generating manually 237
 Java-supplied 311
 subclasses 235
 superclasses 214
 transient fields 224
 transient versions 203
 using as transient 319
 persistent objects
 associated session 28
 definition 6
 destroying 142
 evicting all 136
 exporting 80
 external references 109
 garbage collection 61
 hollow after abort() 141
 hollow after commit() 126
 hollow after evict() 134
 identity 119
 is this object persistent? 103
 migrating 78
 multiple representations 34
 nonexported 131
 object state, specifying 119
 optimizing retrieval 230
 readable after abort() 141
 readable after commit() 127
 retrieving 102
 serializing 317
 specifying UPDATE transaction type 86
 stale after abort() 140
 stale after commit() 125
 stale after evict() 134
 transient after abort() 141
 transient after commit() 130
 transient after evictAll() 136
 transient fields 146
 updatable after abort() 141
 updatable after commit() 129

- `Persistent.preDestroyPersistent()`
 - method 142
- `PersistentTypeSummary` class 273
- `Placement.getSession()` method 27
- planning for clusters 78
- `postInitializeContents()` hook
 - method 227
- postprocessor
 - annotated class files, location 215
 - annotated class files, managing 208
 - annotated classes, subclasses 235
 - applications, complex 210
 - batches 198
 - Class could not be found error 202
 - CLASSPATH requirements 200
 - command line, sample 201
 - consistency 213
 - conventions 252
 - customizing 226
 - debugger 217
 - destination directory background 202
 - destination directory requirement 201
 - duplicate file specifications 205
 - errors and warnings 215
 - example of multiple persistence
 - modes 205
 - example of running 202
 - file name interpretation 203
 - file not found 206
 - final fields 215
 - hollow object constructor 229
 - hook methods, sample 227
 - how it works 213
 - input file 233
 - introduction 196
 - jar files 206
 - Java classes, modifying 327
 - limitations 236
 - new packages 221
 - nonpersistent classes 219
 - nonpersistent methods 234
 - not used by JDD 292
 - objects, optimizing retrieval of 230
 - optimizations 218
 - optimizations can cause problems 148
 - PATH requirements 201
 - persistence mode options 204
 - persistence-aware classes 212
 - preparations 200
 - previously annotated classes 206
 - processing order 204
 - running 200, 323
 - static fields 216
 - superclasses, modifications 214
 - testing 232
 - transient classes 223
 - transient fields 224
 - translatepackage option 221
 - zip files 206
- `preClearContents()` hook method 227
- `preFlushContents()` hook method 227
- primary objects 10
- product property 42
- properties
 - `com.odi.addressSpaceSize` 38
 - `com.odi.applicationName` 39
 - `com.odi.cachedObjectCount` 39
 - `com.odi.cachedObjectTransactionAge` 39
 - `com.odi.cacheSize` 40
 - `com.odi.disableObjectCaching` 40
 - `com.odi.disableWeakReferences` 41
 - `com.odi.gc.displayGarbage` 64
 - `com.odi.gc.lockTimeOut` 63
 - `com.odi.gc.reachableObjects` 64
 - `com.odi.gc.reclaimedObjects` 64
 - `com.odi.gc.retries` 63
 - `com.odi.gc.retryInterval` 63
 - `com.odi.gc.transactionPriority` 63
 - `com.odi.migrateUnexportedStrings` 42
 - `com.odi.ObjectStoreLibrary` 42

- `com.odi.password` 42
- `com.odi.product` 42
- `com.odi.stringPoolSize` 45
- `com.odi.useImmediateStrings` 46
- `com.odi.user` 42
- garbage collection 63
- parameter 38
- system 37
- `trapUnregisteredType` 46

PSE Pro

- code example of use with OSJI 44
- copying data to ObjectStore 42

PSEPro

- using with OSJI 42

Q

queries

- creating 173
- example 174
- executing 182
- expression syntax in JDD 296
- free variables 181
- in JDD 296
- indexes 184
- introduction 173
- limitations 183
- maps 166
- multistep indexes 186
- null values 183
- operators 177
- optimizing for indexes 190
- optimizing in JDD 306
- pattern matching 179
- sample program 175
- string literals 178
- syntax 177
- unsupported 178
- `-quiet` option 232, 329
- `-quietclass` option 215, 330
 - suppressing warnings 232
- `-quietfield` option 215, 330

- suppressing warnings 232

R

- reachability 101
- read locks 37
- reading notifications 287
- read-only
 - database open type 57
- receiving notifications 284
- recovery 3
- references
 - checking 338
 - cross-segment 77
 - destroying sources 142
 - destroying targets 145
 - external 109
 - from evicted objects 133
 - to transient instances 224
- reflection API 207
- registering classes
 - manual annotation 238
 - postprocessor 196
- relationships, JDD
 - example 303
 - linked objects 302
 - many-to-many 301
 - one-to-many 300
 - one-to-one 300
 - types 299
- remote machines 2
- remote method invocation
 - See* RMI
- `RestartableAbortException` 93
- restarting aborted transactions 94
- RETAIN constants
 - See* `ObjectStore.RETAIN`
- retaining objects
 - abort transaction 138
 - commit read-only 127
 - commit transient 130
 - commit update 129

- commit(ObjectStore.RETAIN
HOLLOW) 126
- default abort retain state 139
- default checkpoint retain state 266
- default commit retain state 124
- evict active 135
- evict hollow 134
- eviction 132
- nonexported 131
- retain argument 122
- retrieving notifications 286
- RMI
 - preparing to serialize 128
 - serializing for 317
 - using persistence-capable classes 319

S

- saving modifications
 - committing transactions 122
 - evicting objects 132
- schema evolution
 - needed when 66
 - not required for JDD 292
 - preparation 67
 - procedure 65
 - serialization sample code 70
 - superclasses 68
- schema information 272
 - installing 271
- schema option 333
- schema segment 36
- security
 - notifications 282
- Segment.acquireLock() method 268
- Segment.createCluster() method 54
- Segment.GC() method 63
- Segment.getCluster() method 54
- Segment.getObjects() method 108
- Segment.setDefaultCluster()
 - method 55
- segments

- description 53
- garbage collection 63, 78
- listing 54
- locking 267
- objects, iterating through 108
- objects, wrong placement of 81
- references across segments 77
- situations to avoid 101
- storing objects in particular ones 101
- transient 56
- sending notifications 285
- serializability 262
- serialization
 - persistent objects 317
 - sample code for schema evolution 70
- server
 - description 4
- ServerRefusedConnectionException 93
- ServerRestartedException 93
- Session.createGlobal() method 23
- Session.getCurrent() method 27, 33
- Session.getGlobal() method 24, 27
- Session.getName() method 25
- Session.isActive() method 27
- Session.join() method 30
- Session.leave() method 32
- Session.of(object) method 27
- Session.ofThread(thread) method 27
- Session.terminate() method 26
- sessions
 - associated objects 28
 - calls that do not imply 29
 - calls that imply 29
 - creating 23
 - definition 20
 - global session 23
 - is one active? 27
 - join rules 29
 - metaobjects 36
 - names 25
 - nonglobal 24

- objects, copies of 21
- obtaining 26
- properties, specifying 37
- shutting down 26
- threads 27
- threads not associated 33
- threads, explicitly adding 30
- threads, relationship to 27
- threads, removing 32
- transactions 25
- `SetLazyWriteLocking()` method 37
- sharing data between PSE Pro and
 - `ObjectStore` 42
- `-showData` option 340
- `-showObjs` option 339
- shutting down sessions 26
- 64 bit platforms
 - transient C++ peer objects 54
- `Cluster.of()` method
 - transient C++ peer objects on 64 bit
 - platforms 54
- transient C++ peer objects
 - 64 bit platforms 54
- stale persistent objects
 - aborting transactions 140
 - attempts to access 146
 - committing transactions 125
 - definition 8
 - `evict()` 134
- static fields
 - postprocessor handling 216
 - session ownership 35
- `Step into` command 218
- `Step out` command 218
- `Step over` command 218
- storing external references 110
- storing objects
 - has this object been stored? 103
 - in particular clusters 101
 - in particular segments 101
 - persistence 100

- procedure 100
- storing objects in clusters 55
- string pool size 45
- strings
 - destroying 316
 - embedded 325
 - garbage collection 62
 - making them persistent 315
 - pool size 45
 - queries 178
 - unexported, migrating 42
- `stublib.jar` file 319
- subscribing to receive notifications 284
- `-summary` option
 - description 330
 - how to use 273
- superclasses
 - abstract 252
 - modifications for persistence 214
 - persistence-aware classes 213
 - schema evolution 68
- superindexes in JDD 306
- `-suppress` option 336
- suppressing generation of methods 336
- synchronization 124
- `-synchronize` option 336
- `-sysclasspath` option 324
- system crash 3
- system properties 37

T

- terminating sessions 26
- testing postprocessor 232
- third-party collections 194
- threads
 - already initialized? 33
 - committing a transaction, effect of 35
 - cooperating 31
 - joined to session? 33
 - joining session explicitly 30
 - noncooperating 32, 33

- not joined to session 33
- notifications 281
- objects, evicting 136
- persistent objects, access to 34
- removing from session 32
- sessions 27
- synchronizing 32
- transaction boundaries 97
- tombstones
 - destroyed objects 142
 - persistent garbage collection 62
- tools.jar file 17
- Transaction class
 - description 86
- Transaction.abort()
 - canceling modifications 138
 - example 140
 - general discussion 91
 - retain 92
- Transaction.abort(retain)
 - specifying object state 138
- Transaction.abort(RETAIN_
 - HOLLOW) 141
- Transaction.abort(RETAIN_
 - READONLY) 141
- Transaction.abort(RETAIN_STALE) 140
- Transaction.abort(RETAIN_
 - UPDATE) 141
- Transaction.begin() method 86
- Transaction.commit()
 - general discussion 90
 - saving modifications 123
- Transaction.commit(retain)
 - general discussion 90
 - specifying object state 123
- Transaction.current() method 88
- Transaction.getPriority() method 271
- Transaction.getSession(thread)
 - method 27
- Transaction.setDefaultAbortRetain()
 - method 139
- Transaction.setDefaultCommitRetain(
 -) method 124, 266
- Transaction.setDefaultRetain()
 - method 124, 139, 266
- Transaction.setPriority() method 271
- transactions
 - aborted, retrying 94
 - aborting 91, 92
 - aborting to cancel changes 138
 - boundaries, determining 96
 - checkpoint 264
 - committing
 - description 90
 - setting object state 122
 - situations to avoid 101
 - deadlocks 92
 - ending 89
 - evicting objects outside 137
 - length 260
 - multiversion concurrency control 261
 - nested 88
 - notifications 281
 - priority 271
 - RETAIN_HOLLOW 126
 - RETAIN_READONLY 127
 - RETAIN_STALE 125
 - RETAIN_TRANSIENT 130
 - RETAIN_UPDATE 129
 - sessions 25
 - starting 86
 - transaction object, obtaining 88
 - update and read-only, comparison 87
- transient and persistence-capable versions
 - of same class 223
- transient database 56
- transient fields
 - annotations, manual 248
 - annotations, preventing 231
 - initialization 225
 - persistence-capable classes,
 - behavior 146
 - postprocessor 224

- transient instance of persistence-capable class 224
- transient objects 10
- transient segment 56
- transient version of class file 203
- transient views of collections 166
- transientfield option 225, 330
- transitive persistence
 - becoming persistent 100
 - definition 11
- translatepackage option 221, 331
- trapping unregistered types 46
- trapUnregisteredType property 46
- troubleshooting
 - access not allowed 218
 - authentication required 42
 - bad command or file name 217, 322
 - class could not be found 202
 - ClassCastException 46, 149
 - destroyed objects, references to 145
 - OutOfMemoryError
 - postprocessor 206
 - storing large objects 154
 - retaining for read or update 128
 - trapping unregistered types 46
 - UnregisteredTypeException 149
- two sessions
 - static variables 35
 - two object copies 21
- Type.addToExtent() method in JDD 296
- Type.entities() method in JDD 294
- Type.extent() method in JDD 294
- TypeQuery class in JDD 297
- types, JDD
 - adding indexes 295
 - defining 294
 - defining attributes in 295
 - finding in database 296
 - overview 292
 - querying 296

U

- unexported strings property 42
- Unicode strings 313
- unknown types 149
- unregistered type property 46
- UnregisteredType class 149
- UnregisteredTypeException 149
- update
 - database open type 57
- updating objects 118
- upgrading databases 73
- upgrading PSE Pro to ObjectStore 42
- user property 42
- UTF8 encoding 313
- utilities
 - garbage collection 64
 - oscompact 64
 - osdump 73
 - osjcheckdb 77, 338
 - osjshowdb 77, 339
 - osjversion 341
 - osload 73

V

- variable initializers 249
- verbose option 232, 331
- version information 341
- very large aggregations 161
- views of maps 166

W

- weak references 41, 122
- weak references property 41
- wrapper classes
 - identity 312
 - persistence-capable 311
 - queries 178
- write locks 37

Z

zip files 206

