

IBM WebSphere eXtreme Scale Version 7.1

Programming Guide

June 15, 2011



This edition applies to version 7, release 1, of WebSphere eXtreme Scale and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2009, 2011.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures v

Tables vii

About the *Programming Guide* ix

Chapter 1. Planning to develop applications 1

WebSphere eXtreme Scale programming interfaces. 1
Class loader and classpath considerations. 2
Directory conventions 2
Setting up a stand-alone development environment 3
Running a WebSphere eXtreme Scale client or server application with Apache Tomcat in Rational Application Developer 5
Running an integrated client or server application with WebSphere Application Server in Rational Application Developer 8

Chapter 2. Accessing data with client applications 9

ObjectGrid interface 9
BackingMap interface 12
Connecting to a distributed ObjectGrid 16
Interacting with an ObjectGrid using ObjectGridManager. 17
 createObjectGrid methods 17
 getObjectGrid methods 21
 removeObjectGrid methods 21
 Controlling the life cycle of an ObjectGrid 22
 Accessing the ObjectGrid shard. 24
Data access with indexes (Index API). 25
 Indexing 28
 Plug-ins for custom indexing of cache objects 30
 Plug-ins for indexing data 33
Accessing data in WebSphere eXtreme Scale 33
 CopyMode best practices. 36
 Byte array maps. 41
Using Sessions to access data in the grid 43
 SessionHandle for routing 47
 SessionHandle integration 47
Caching objects with no relationships involved (ObjectMap API). 50
 Introduction to ObjectMap 50
 Dynamic maps 53
 ObjectMap and JavaMap 56
 Maps as FIFO queues 57
Caching objects and their relationships (EntityManager API) 60
 Defining an entity schema 61
 EntityManager in a distributed environment 71
 Interacting with EntityManager. 74
 EntityManager fetch plan support 81
 EntityManager interface performance impact 84

Entity query queues 87
EntityTransaction interface 91
Entity manager tutorial: Overview. 91
Retrieving entities and objects (Query API). 92
 Querying data in multiple time zones 95
 Inserting data for different time zones 97
 Using the ObjectQuery API 97
 EntityManager Query API 102
 Reference for eXtreme Scale queries 106
 Query performance tuning 116
 Using objects other than keys to find partitions (PartitionableKey interface). 127
Programming for transactions 128
 Transaction processing overview 128
 Using locking 141
 Transaction isolation 152
 Optimistic collision exception 154
 Running parallel business logic on the data grid (DataGrid API) 155
Configuring clients with WebSphere eXtreme Scale 160
 Tracking map updates by an application 165
 Enabling client-side map replication. 168

Chapter 3. Accessing data with the REST data service 169

Operations with the REST data service 169
Request protocols for the REST data service 171
 Retrieve requests with the REST data service 171
 Retrieving non-entities with REST data services 178
 Insert requests with REST data services 183
 Update requests with REST data services 187
 Delete requests with REST data services 191
Optimistic concurrency 192

Chapter 4. Programming with system APIs and plug-ins 193

Introduction to plug-ins 193
Plug-ins for evicting cache objects 195
 TimeToLive (TTL) evictor 198
 Plug in a pluggable evictor. 200
 Writing a custom evictor 203
 Plug-in evictor performance best practices. 207
Plug-ins for transforming cached objects 209
 Developing custom arbiters for multi-master replication 209
 ObjectTransformer plug-in 210
Plug-ins for versioning and comparing cache objects. 216
Plug-ins for providing event listeners 221
 MapEventListener plug-in 222
 ObjectGridEventListener plug-in 223
Plug-ins for custom indexing of cache objects. 225
 Using a composite index 227
Plug-ins for communicating with persistent stores 230
 Writing a loader 232

| | |
|---|-----|
| JPA loader programming considerations | 236 |
| JPAEntityLoader plug-in. | 238 |
| Using a loader with entity maps and tuples | 240 |
| Writing a loader with a replica preload controller. | 245 |
| Plug-ins for managing transaction life cycle events | 250 |
| Transaction processing overview | 254 |
| Introduction to plug-in slots | 254 |
| External transaction managers. | 256 |
| WebSphereTransactionCallback plug-in. | 259 |

Chapter 5. Programming for administrative tasks 261

| | |
|---|-----|
| Embedded server API | 261 |
| Using the embedded server API to start and stop servers | 263 |
| Monitoring with the statistics API | 265 |
| Monitoring with WebSphere Application Server PMI | 268 |
| Enabling PMI | 268 |
| Retrieving PMI statistics. | 270 |
| PMI modules | 272 |

Chapter 6. Programming for JPA integration 281

| | |
|---|-----|
| JPA Loaders | 281 |
| Developing client-based JPA loaders. | 283 |
| Client-based JPA preload utility overview | 284 |
| Example: Preloading a map with the ClientLoader interface | 285 |
| Example: Reloading a map with the ClientLoader interface | 286 |
| Example: Calling a client loader | 286 |
| Example: Creating a custom client-based JPA loader | 287 |
| Developing a client-based JPA loader with a DataGrid agent. | 287 |
| Starting the JPA time-based updater. | 291 |
| JPA time-based data updater | 294 |

Chapter 7. Programming for Spring integration 297

| | |
|---------------------------|-----|
| Spring framework. | 297 |
|---------------------------|-----|

| | |
|--|-----|
| Managing transactions with Spring | 298 |
| Spring managed extension beans | 300 |
| Spring extension beans and namespace support | 300 |
| Starting a container server with Spring. | 303 |

Chapter 8. Programming for security 307

| | |
|---|-----|
| Security API. | 307 |
| Client authentication programming | 308 |
| Client authorization programming | 326 |
| Data grid authentication. | 333 |
| Local security | 333 |

Chapter 9. Performance considerations for application developers 339

| | |
|--|-----|
| JVM tuning for WebSphere eXtreme Scale | 339 |
| CopyMode best practices | 341 |
| Byte array maps | 346 |
| Plug-in evictor performance best practices. | 348 |
| Locking performance best practices | 350 |
| Serialization performance | 351 |
| ObjectTransformer interface best practices | 353 |
| Query performance tuning | 354 |

Chapter 10. Troubleshooting 357

| | |
|---|-----|
| Enabling logging | 357 |
| Collecting trace. | 358 |
| Trace options | 359 |
| Troubleshooting client connectivity | 361 |
| Troubleshooting loaders | 361 |
| Troubleshooting deadlocks | 362 |
| Troubleshooting security. | 367 |
| IBM Support Assistant for WebSphere eXtreme Scale | 368 |
| Messages. | 369 |
| Release notes | 369 |

Notices 371

Trademarks 373

Index 375

Figures

1. The interaction of the query with the ObjectGrid object maps and how a schema is defined for classes and associated with an ObjectGrid map 98
2. The interaction of the query with the ObjectGrid object maps and how the entity schema is defined and associated with an ObjectGrid map. 103
3. Loader. 231
4. ObjectGridModule module structure. . . . 272
5. ObjectGridModule module structure example 273
6. mapModule structure. 274
7. mapModule module structure example 274
8. hashIndexModule module structure 276
9. hashIndexModule module structure example 276
10. agentManagerModule structure 277
11. agentManagerModule structure example 278
12. queryModule structure 279
13. QueryStats.jpg queryModule structure example 279
14. JPA Loader architecture 282
15. Client loader that uses JPA implementation to load the ObjectGrid 284
16. Periodic refresh 294
17. Flow of client authentication and authorization 307

Tables

| | | | |
|--|-----|---|-----|
| 1. ObjectGrid interface | 9 | 8. Permissions to a server-hosted ObjectMap | 328 |
| 2. Other methods. | 95 | 9. Single key deadlocks scenario | 364 |
| 3. Key to BNF summary. | 114 | 10. Single key deadlocks, continued | 364 |
| 4. Lock mode compatibility matrix | 142 | 11. Single key deadlocks, continued | 364 |
| 5. Client loader modes | 285 | 12. Single key deadlocks, continued | 365 |
| 6. List of methods and the required MapPermission | 327 | 13. Ordered multiple key deadlock scenario | 365 |
| 7. List of methods and the required ObjectGridPermission. | 328 | 14. Ordered multiple key deadlock scenario, continued. | 366 |
| | | 15. Out of order with U lock scenario. | 366 |

About the *Programming Guide*

The WebSphere® eXtreme Scale documentation set includes three volumes that provide the information necessary to use, program for, and administer the WebSphere eXtreme Scale product.

WebSphere eXtreme Scale library

The WebSphere eXtreme Scale library contains the following books:

- The *Product Overview* contains a high-level view of WebSphere eXtreme Scale concepts, including use case scenarios, and tutorials.
- The *Installation Guide* describes how to install common topologies of WebSphere eXtreme Scale.
- The *Administration Guide* contains the information necessary for system administrators, including how to plan application deployments, plan for capacity, install and configure the product, start and stop servers, monitor the environment, and secure the environment.
- The *Programming Guide* contains information for application developers on how to develop applications for WebSphere eXtreme Scale using the included API information.

To download the books, go to the WebSphere eXtreme Scale library page.

You can also access the same information in this library in the WebSphere eXtreme Scale information center.

Who should use this book

This book is intended primarily for application developers.

Getting updates to this book

You can get updates to this book by downloading the most recent version from the WebSphere eXtreme Scale library page.

How to send your comments

Contact the documentation team. Did you find what you needed? Was it accurate and complete? Send your comments about this documentation by e-mail to wasdoc@us.ibm.com.

Chapter 1. Planning to develop applications

Set up your development environment and learn where to find details about available programming interfaces.

WebSphere eXtreme Scale programming interfaces

WebSphere eXtreme Scale provides several features that are accessed programmatically using the Java programming language through application programming interfaces (APIs) and system programming interfaces.

WebSphere eXtreme Scale APIs

When you are using eXtreme Scale APIs, you must distinguish between transactional and non-transactional operations. A transactional operation is an operation that is performed within a transaction. ObjectMap, EntityManager, Query, and DataGrid API are transactional APIs that are contained inside the Session that is a transactional container. Non-transactional operations have nothing to do with a transaction, such as configuration operations.

The ObjectGrid, BackingMap, and plug-in APIs are non-transactional. The ObjectGrid, BackingMap, and other configuration APIs are categorized as ObjectGrid Core API. Plug-ins are for customizing the cache to achieve the functions that you want, and are categorized as the System Programming API. A plug-in in eXtreme Scale is a component that provides a certain type of function to the pluggable eXtreme Scale components that include ObjectGrid and BackingMap. A feature represents a specific function or characteristic of an eXtreme Scale component, including ObjectGrid, Session, BackingMap, ObjectMap, and so on. Typically, features are configurable with configuration APIs. Plug-ins can be built-in, but might require that you develop your own plug-ins in some situations.

You can normally configure the ObjectGrid and BackingMap to meet your application requirements. When the application has special requirements, consider using specialized plug-ins. WebSphere eXtreme Scale might have built-in plug-ins that meet your requirements. For example, if you need a peer-to-peer replication model between two local ObjectGrid instances or two distributed eXtreme Scale grids, the built-in JMSObjectGridEventListener is available. If none of the built-in plug-ins can solve your business problems, refer to the System Programming API to provide your own plug-ins.

ObjectMap is a simple map-based API. If the cached objects are simple and no relationship is involved, the ObjectMap API is ideal for your application. If object relationships are involved, use the EntityManager API, which supports graph-like relationships.

Query is a powerful mechanism for finding data in the ObjectGrid. Both Session and EntityManager provide the traditional query capability.

The DataGrid API is a powerful computing capability in a distributed eXtreme Scale environment that involves many machines, replicas, and partitions. Applications can run business logic in parallel in all of the nodes in the distributed eXtreme Scale environment. The application can obtain the DataGrid API through the ObjectMap API.

7.0.0.0 FIX 2+ The WebSphere eXtreme Scale REST data service is a Java HTTP service that is compatible with Microsoft WCF Data Services (formally ADO.NET Data Services) and implements the Open Data Protocol (OData). The REST data service allows any HTTP client to access an eXtreme Scale grid. It is compatible with the WCF Data Services support that is supplied with the Microsoft .NET Framework 3.5 SP1. RESTful applications can be developed with the rich tooling provided by Microsoft Visual Studio 2008 SP1. For more details, refer to the eXtreme Scale REST data service user guide.

Class loader and classpath considerations

Since eXtreme Scale stores Java objects in the cache by default, you must define classes on the classpath wherever the data is accessed.

Specifically, eXtreme Scale client and container processes must include the classes or JARs in the classpath when starting the process. When designing an application for use with eXtreme Scale, separate out any business logic from the persistent data objects.

See Class loading in the WebSphere Application Server Network Deployment information center for more information.

For considerations within a Spring Framework setting, see the packaging section under the topic on integrating with Spring framework in the *Programming Guide*.

For settings related to using the WebSphere eXtreme Scale instrumentation agent, see the instrumentation agent topic in the *Programming Guide*.

Directory conventions

The following directory conventions are used throughout the documentation to must reference special directories such as `wxs_install_root` and `wxs_home`. You access these directories during several different scenarios, including during installation and use of command-line tools.

`wxs_install_root`

The `wxs_install_root` directory is the root directory where WebSphere eXtreme Scale product files are installed. The `wxs_install_root` directory can be the directory in which the trial zip file is extracted or the directory in which the WebSphere eXtreme Scale product is installed.

- Example when extracting the trial:

Example: `/opt/IBM/WebSphere/eXtremeScale`

- Example when WebSphere eXtreme Scale is installed to a stand-alone directory:

Example: `/opt/IBM/eXtremeScale`

- Example when WebSphere eXtreme Scale is integrated with WebSphere Application Server:

Example: `/opt/IBM/WebSphere/AppServer`

`wxs_home`

The `wxs_home` directory is the root directory of the WebSphere eXtreme Scale product libraries, samples and components. This is the same as the `wxs_install_root` directory when the trial is extracted. For stand-alone installations, the `wxs_home` directory is the ObjectGrid sub-directory within the `wxs_install_root` directory. For installations that are integrated with

WebSphere Application Server, this directory is the optionalLibraries/ObjectGrid directory within the *wxs_install_root* directory.

- Example when extracting the trial:
Example: /opt/IBM/WebSphere/eXtremeScale
- Example when WebSphere eXtreme Scale is installed to a stand-alone directory:
Example: /opt/IBM/eXtremeScale/ObjectGrid
- Example when WebSphere eXtreme Scale is integrated with WebSphere Application Server:
Example: /opt/IBM/WebSphere/AppServer/optionalLibraries/ObjectGrid

was_root

The *was_root* directory is the root directory of a WebSphere Application Server installation:

Example: /opt/IBM/WebSphere/AppServer

restservice_home

The *restservice_home* directory is the directory in which the WebSphere eXtreme Scale REST data service libraries and samples are located. This directory is named *restservice* and is a sub-directory under the *wxs_home* directory.

- Example for stand-alone deployments:
Example: /opt/IBM/WebSphere/eXtremeScale/ObjectGrid/restservice
- Example for WebSphere Application Server integrated deployments:
Example: /opt/IBM/WebSphere/AppServer/optionalLibraries/ObjectGrid/restservice

tomcat_root

The *tomcat_root* is the root directory of the Apache Tomcat installation.

Example: /opt/tomcat5.5

wasce_root

The *wasce_root* is the root directory of the WebSphere Application Server Community Edition installation.

Example: /opt/IBM/WebSphere/AppServerCE

java_home

The *java_home* is the root directory of a Java Runtime Environment (JRE) installation.

Example: /opt/IBM/WebSphere/eXtremeScale/java

samples_home

The *samples_home* is the directory in which you extract the sample files that are used for tutorials.

Example: /wxs-samples/

Setting up a stand-alone development environment

Configure an Eclipse-based integrated development environment to build and run a Java SE application with the stand-alone version of WebSphere eXtreme Scale.

Before you begin

Install the WebSphere eXtreme Scale product into a new or empty directory and apply the latest WebSphere eXtreme Scale cumulative fix pack. You can also use the WebSphere eXtreme Scale trial version by unzipping the zip file. For more information on installation, see the information on installing the stand-alone WebSphere eXtreme Scale or WebSphere eXtreme Scale Client in the *Administration Guide*.

Procedure

- Configure Eclipse to build and run a Java SE application with WebSphere eXtreme Scale.
 1. Define a user library to allow your application to reference WebSphere eXtreme Scale application programming interfaces.
 - a. In your Eclipse or IBM® Rational® Application Developer environment, click **Window > Preferences**.
 - b. Expand the **Java > Build Path** branch and select **User Libraries**. Click **New**.
 - c. Select the eXtreme Scale user library. Click **Add JARs**.
 - 1) Browse and select the `objectgrid.jar` or `ogclient.jar` files from the `wxs_root/lib` directory. Click **OK**. Select the `ogclient.jar` file if you are developing client applications or local, in-memory caches. If you are developing and testing eXtreme Scale servers, use the `objectgrid.jar` file.
 - 2) To include Javadoc for the ObjectGrid APIs, select the Javadoc location for the `objectgrid.jar` or `ogclient.jar` file that you added in the previous step. Click **Edit**. In the Javadoc location path box, type the following web address:
`http://www.ibm.com/developerworks/wikis/extremescale/docs/api/`
 - d. Click **OK** to apply the settings and close the Preferences window.
The eXtreme Scale libraries are now in the build path for the project.
 2. Add the user library to your Java project.
 - a. From the package explorer, right-click the project and select **Properties**.
 - b. Select the **Libraries** tab.
 - c. Click **Add Library**.
 - d. Select **User Library**. Click **Next**.
 - e. Select the eXtreme Scale user library that you configured earlier.
 - f. Click **OK** to apply the changes and close the Properties window.
- Run a Java SE application with eXtreme Scale with Eclipse. Create a run configuration to execute your application.
 1. Configure Eclipse to build and run a Java SE application with eXtreme Scale. From the **Run** menu select **Run Configurations**.
 2. Right-click the Java Application category and select **New**.
 3. Select the new run configuration, named *New_Configuration*.
 4. Configure the profile.
 - **Project** (on main tabbed page): *your_project_name*
 - **Main Class** (on main tabbed page): *your_main_class*
 - **VM arguments** (on arguments tabbed page):
`-Djava.endorsed.dirs=wxs_root/lib/endorsed`

Problems with the **VM Arguments** often occur because the path to `java.endorsed.dirs` must be an absolute path with no variables or shortcuts. Other common setup problems involve the Object Request Broker (ORB). You might see the following error. Refer to *Configuring a custom Object Request Broker* for more information:

```
Caused by: java.lang.RuntimeException: The ORB that comes
with the Sun Java implementation does not work with
ObjectGrid at this time.
```

If you do not have the `objectGrid.xml` or `deployment.xml` accessible to the application, you might see the following error:

```
Exception in thread "P=211046:0=0:CT" com.ibm.websphere.objectgrid.
ObjectGridRuntimeException: Cannot start OG container at
Client.startTestServer(Client.java:161) at Client.
main(Client.java:82) Caused by: java.lang.IllegalArgumentException:
The objectGridXML must not be null at com.ibm.websphere.objectgrid.
deployment.DeploymentPolicyFactory.createDeploymentPolicy
(DeploymentPolicyFactory.java:55) at Client.startTestServer(Client.
java:154) .. 1 more
```

5. Click **Apply** and close the window, or click **Run**.

Running a WebSphere eXtreme Scale client or server application with Apache Tomcat in Rational Application Developer

Whether you have a client or server application, use the same basic steps to run the application in Apache Tomcat in Rational Application Developer. For a client application, you want to configure and run a web application to use a WebSphere eXtreme Scale client in Rational Application Developer. Follow these instructions to create a web project for running a WebSphere eXtreme Scale catalog service or container. For a server application, you want to enable a Java EE application in Rational Application Developer interface with a stand-alone installation of WebSphere eXtreme Scale. Follow these instructions to configure a Java EE application project for using the WebSphere eXtreme Scale client library.

Before you begin

Install the WebSphere eXtreme Scale Trial or full product.

- Install the stand-alone version of the WebSphere eXtreme Scale Version 7.1 product.
- Download and extract the WebSphere eXtreme Scale trial version.
- Install Apache Tomcat Version 6.0 or later.
- Install Rational Application Developer and create a Java EE web application.

Procedure

1. Add WebSphere eXtreme Scale runtime library to your Java EE build path.
Client application In this scenario, you want to configure and run a web application to use a WebSphere eXtreme Scale client in Rational Application Developer.
 - a. **Window > Preferences > Java > Build Path > User Libraries**. Click **New**.
 - b. Enter a **User library name** of `eXtremeScaleClient`, and click **OK**.
 - c. Click **Add Jars...**, and navigate to and select the `wxs_home/lib/ogclient.jar` file. Click **Open**.

- d. Optional: (Optional) To add Javadoc, select Javadoc location and click **Edit...** In the Javadoc location path, you can either enter the URL of the API documentation, or you can download the API documentation.
 - To use the online API documentation, enter `http://www.ibm.com/developerworks/wikis/extremescale/docs/api/` in the Javadoc location path.
 - To download the API documentation, go to the WebSphere eXtreme Scale API documentation download page. For the Javadoc location path, enter your local download location.
- e. Click **OK**.
- f. Click **OK** to close out the User Libraries dialogue.
- g. Click **Project > Properties**.
- h. Click **Java Build Path**.
- i. Click **Add Library**.
- j. Select **User Library**. Click **Next**.
- k. Check the **eXtremeScaleClient** library and click **Finish**.
- l. Click **OK** to close the **Project Properties** dialog.

Server application In this scenario, you want to configure and run a web application to run an embedded WebSphere eXtreme Scale server in Rational Application Developer.

- a. Click **Window > Preferences > Java > Build Path > User Libraries**. Click **New**.
 - b. Enter a **User library name** of `eXtremeScale`, and click **OK**.
 - c. Click **Add Jars...**, and select `wxs_home/lib/objectgrid.jar`. Click **Open**.
 - d. (Optional) To add Javadoc, select Javadoc location and click **Edit...** In the Javadoc location path, Enter `http://www.ibm.com/developerworks/wikis/extremescale/docs/api/`.
 - e. Click **OK**.
 - f. Click **OK** to close out the User Libraries dialogue.
 - g. Click **Project > Properties**.
 - h. Click **Java Build Path**.
 - i. Click **Add Library**.
 - j. Select **User Library**. Click **Next**.
 - k. Check the **eXtremeScaleClient** library and click **Finish**.
 - l. Click **OK** to close the **Project Properties** dialog.
2. Define Tomcat Server for our project.
 - a. Ensure that you are in the J2EE perspective and click the **Servers** tab in the bottom pane. You can also click **Window > Show View > Servers**.
 - b. Right-click in the Servers pane, and choose **New > Server**.
 - c. Choose **Apache, Tomcat v6.0 Server**. Click **Next**.
 - d. Click **Browse...** Select `tomcat_root`. Click **OK**.
 - e. Click **Next**.
 - f. Select your Java EE application in the left Available pane and click **Add >** to move it to the right Configured pane on the server, and click **Finish**.
 3. Resolve any remaining errors for the Project. Use the following steps to eliminate errors in the Problems pane:
 - a. Click **Project > Clean > project_name**. Click **OK**. Build the project.

- b. Right-click on the Java EE project, and choose **Build Path > Configure Build Path**.
- c. Click the **Libraries** tab. Ensure that the path is configured properly:
 - **For client applications:** Ensure that Apache Tomcat, eXtremeScaleClient, and Java 1.5 JRE are on the path.
 - **For server applications:** Ensure that Apache Tomcat, eXtremeScale, and Java 1.5 JRE are on the path.
4. Create a run configuration to run your application.
 - a. From the **Run** menu, select **Run Configurations**.
 - b. Right-click the Java Application category and select **New**.
 - c. Select the new run configuration, named *New_Configuration*.
 - d. Configure the profile.
 - **Project** (on main tabbed page): *your_project_name*
 - **Main Class** (on main tabbed page): *your_main_class*
 - **VM arguments** (on arguments tabbed page):
-Djava.endorsed.dirs=wxs_root/lib/endorsed

Problems with the **VM Arguments** often occur because the path to `java.endorsed.dirs` must be an absolute path with no variables or shortcuts.

Other common setup problems involve the Object Request Broker (ORB). You might see the following error. See *Configuring a custom Object Request Broker* for more information:

Caused by: `java.lang.RuntimeException: The ORB that comes with the Sun Java implementation does not work with ObjectGrid at this time.`

If you do not have the `objectGrid.xml` or `deployment.xml` files accessible to the application, you might see the following error:

```
Exception in thread "P=211046:0=0:CT" com.ibm.websphere.objectgrid.ObjectGridRuntimeException:
Cannot start OG container
    at Client.startTestServer(Client.java:161)
    at Client.main(Client.java:82)
Caused by: java.lang.IllegalArgumentException: The objectGridXML must not be null
    at com.ibm.websphere.objectgrid.deployment.DeploymentPolicyFactory.createDeploymentPolicy
      (DeploymentPolicyFactory.java:55)
    at Client.startTestServer(Client.java:154)
... 1 more
```

5. Click **Apply** and close the window, or click **Run**.

What to do next

After you configure and run a web application with WebSphere eXtreme Scale client in Rational Application Developer, you can develop a servlet. This servlet uses the WebSphere eXtreme Scale APIs to store and retrieve data from a remote data grid.

After you enable a Java EE application in Rational Application Developer interface with a stand-alone installation of WebSphere eXtreme Scale, you can develop a servlet that uses the WebSphere eXtreme Scale system APIs to start and stop catalog services.

Running an integrated client or server application with WebSphere Application Server in Rational Application Developer

Configure and run a Java EE application with a WebSphere eXtreme Scale client or server with the WebSphere Application Server runtime embedded in Rational Application Developer. If you are configuring a server, starting WebSphere Application Server automatically starts WebSphere eXtreme Scale.

Before you begin

The following steps are for WebSphere Application Server Version 7.0 with Rational Application Developer Version 7.5. The following steps might vary if you are using different versions of these products.

Install Rational Application Developer with WebSphere Application Server Test Environment extensions.

Install WebSphere eXtreme Scale client or server into the WebSphere Application Server, Version 7.0 Test Environment in the *rad_home\runtimes\base_v7* directory.

Procedure

1. Define eXtreme Scale server that is integrated with WebSphere Application Server for your project.
 - a. In the J2EE perspective, click **Window > Show View > Servers**.
 - b. Right-click in the **Servers** pane. Choose **New > Server**.
 - c. Choose **IBM WebSphere Application Server v7.0**. Click **Next**.
 - d. Select a profile to use. The default is *was70profile1*.
 - e. Enter the server name. The default is *server1*.
 - f. Click **Next**.
 - g. Select your Java EE application in the **Available** pane. Click **Add >** to move it to the **Configured** pane on the server. Click **Finish**.
2. To run the Java EE application, start the application server. Right-click **WebSphere Application Server v7.0** and select **Start**.

Chapter 2. Accessing data with client applications

After you configure your development environment, you can begin to develop applications that create, access, and manage the data in your data grid.

About this task

From the perspective of a client application, using WebSphere eXtreme Scale involves the following main steps:

- Connecting to the catalog service by obtaining a `ClientClusterContext` instance.
- Obtaining a client `ObjectGrid` instance.
- Getting a `Session` instance.
- Getting an `ObjectMap` instance.
- Using the `ObjectMap` methods.

ObjectGrid interface

The following methods allow you to interact with an `ObjectGrid` instance.

Create and initialize

See the `ObjectGridManager` interface topic for the required steps for creating an `ObjectGrid` instance. Two distinct methods exist to create an `ObjectGrid` instance: programmatically or with XML configuration files. See the API documentation for more information.

Get or set and factory methods

Any set methods must be called before you initialize the `ObjectGrid` instance. If you call a set method after the initialize method is called, a `java.lang.IllegalStateException` exception results. Each of the `getSession` methods of the `ObjectGrid` interface also implicitly call the initialize method. Therefore, you must call the set methods before calling any of the `getSession` methods. The only exception to this rule is with the setting, adding, and removing of `ObjectGridEventListener` objects. These objects are allowed to be processed after the initialization processing has completed.

The `ObjectGrid` interface contains the following major methods.

Table 1. ObjectGrid interface. Major methods of ObjectGrid.

| Method | Description |
|---|--|
| <code>BackingMap defineMap(String name);</code> | <code>defineMap</code> : is a factory method to define a uniquely named <code>BackingMap</code> . For more information about backing maps, see <code>BackingMap</code> interface. |
| <code>BackingMap getMap(String name);</code> | <code>getMap</code> : Returns a <code>BackingMap</code> previously defined by calling <code>defineMap</code> . By using this method, you can configure the <code>BackingMap</code> , if it is not already configured through XML configuration. |
| <code>BackingMap createMap(String name);</code> | <code>createMap</code> : Creates a <code>BackingMap</code> , but does not cache it for use by this <code>ObjectGrid</code> . Use this method with the <code>setMaps(List)</code> method of the <code>ObjectGrid</code> interface, which caches <code>BackingMaps</code> for use with this <code>ObjectGrid</code> . Use these methods when you are configuring an <code>ObjectGrid</code> with the Spring Framework. |
| <code>void setMaps(List mapList);</code> | <code>setMaps</code> : Clears any <code>BackingMaps</code> that have been previously defined on this <code>ObjectGrid</code> and replaces them with the list of <code>BackingMaps</code> that is provided. |

Table 1. ObjectGrid interface (continued). Major methods of ObjectGrid.

| Method | Description |
|---|--|
| public Session getSession() throws ObjectGridException, TransactionCallbackException; | getSession: Returns a Session, which provides begin, commit, rollback functionality for a Unit of Work. For more information about Session objects, see Session interface. |
| Session getSession(CredentialGenerator cg); | getSession(CredentialGenerator cg): Get a session with a CredentialGenerator object. This method can only be called by the ObjectGrid client in a client server environment. |
| Session getSession(Subject subject); | getSession(Subject subject): Allows the use of a specific Subject object rather than the one configured on the ObjectGrid to get a Session. |
| void initialize() throws ObjectGridException; | initialize: ObjectGrid is initialized and available for general use. This method is called implicitly when the getSession method is called, if the ObjectGrid is not in an initialized state. |
| void destroy(); | destroy: The framework is disassembled and cannot be used after this method is called. |
| void setTxTimeout(int timeout); | setTxTimeout: Use this method to set the amount of time, in seconds, that a transaction that is started by a session that this ObjectGrid instance created is allowed for completion. If a transaction does not complete within the specified amount of time, the Session that started the transaction is marked as being "timed out". Marking a Session as timed out causes the next ObjectMap method that is invoked by the timed out Session to result in a exception. The Session is marked as rollback only, which causes the transaction to be rolled back even if the application calls the commit method instead of the rollback method after the TransactionTimeoutException exception is caught by the application. A timeout value of 0 indicates that the transaction is allowed unlimited amount of time to complete. The transaction does not time out if a time out value of 0 is used. If this method is not called, then any Session that is returned by the getSession method of this interface has a transaction timeout value set to 0 by default. An application can override the transaction timeout setting on a per Session basis by using the setTransactionTimeout method of the com.ibm.websphere.objectgrid.Session interface. You can also configure transaction timeout with the objectGrid.xml file in the distributed case. |
| int getTxTimeout(); | getTxTimeout: Returns the transaction timeout value in seconds. This method returns the same value that is passed as the timeout parameter on the setTxTimeout method. If the setTxTimeout method was not called, then the method returns 0 to indicate that the transaction is allowed an unlimited amount of time to complete. |
| public int getObjectGridType(); | Returns the type of ObjectGrid. The return value is equivalent to one of the constants declared on this interface: LOCAL, SERVER, or CLIENT. |
| public void registerEntities(Class[] entities); | Register one or more entities based on the class metadata. Entity registration is required prior to ObjectGrid initialization to bind an Entity with a BackingMap and any defined indices. This method may be called multiple times. |
| public void registerEntities(URL entityXML); | Registers one ore more entities from an entity XML file. Entity registration is required prior to ObjectGrid initialization to bind an Entity with a BackingMap and any defined indices. This method may be called multiple times. |
| void setQueryConfig(QueryConfig queryConfig); | Set the QueryConfig object for this ObjectGrid. A QueryConfig object provides query configurations for executing object queries over the maps in this ObjectGrid. |
| //Keywords | |
| void associateKeyword(Serializable parent, Serializable child); | Deprecated: Use Index or query function to get Objects with specific attributes. The associateKeyword method provides a flexible invalidation mechanism based on keywords. This method links the two keywords together in a directional relationship. If parent is invalidated, then the child is also invalidated. Invalidating the child has no impact on the parent. |
| //Security | |
| void setSecurityEnabled() | setSecurityEnabled: Enables security. Security is disabled by default. |
| void setPermissionCheckPeriod(long period); | setPermissionCheckPeriod: This method takes a single parameter that indicates how often to check the permission that is used to allow a client access. If the parameter is 0, all methods ask the authorization mechanism, either JAAS authorization or custom authorization, to check if the current subject has permission. This strategy might cause performance issues depending on the authorization implementation. However, this type of authorization is available if it is required. Alternatively, if the parameter is less than 0, it indicates the number of milliseconds to cache a set of permissions before returning to the authorization mechanism to refresh them. This parameter provides much better performance, but if the backend permissions are changed during this time the ObjectGrid might allow or prevent access even though the backend security provider has been modified. |
| void setAuthorizationMechanism(int authMechanism); | setAuthorizationMechanism: Set the authorization mechanism. The default is SecurityConstants.JAAS_AUTHORIZATION. |

Table 1. ObjectGrid interface (continued). Major methods of ObjectGrid.

| Method | Description |
|--|---|
| setMapAuthorization(MapAuthorization ma); | Deprecated: Use the setObjectGridAuthorization (ObjectGridAuthorization) method instead to plug in custom authorizations. setMapAuthorization: Sets the MapAuthorization plug-in for this ObjectGrid instance. This plug-in can be used to authorize ObjectMap or JavaMap accesses to the principals that are contained in the Subject object. A typical implementation of this plug-in is to retrieve the principals from the Subject object, and then check if the specified permissions are granted to the principals. |
| setSubjectSource(SubjectSource ss); | setSubjectSource: Sets the SubjectSource plugin. This plug-in can be used to get a Subject object that represents the ObjectGrid client. This subject is used for ObjectGrid authorization. The SubjectSource.getSubject method is called by the ObjectGrid runtime when the ObjectGrid.getSession method is used to get a session and the security is enabled. This plug-in is useful for an already authenticated client: it can retrieve the authenticated Subject object and then pass to the ObjectGrid instance. Another authentication is not necessary. |
| setSubjectValidation(SubjectValidation sv); | setSubjectValidation: Sets the SubjectValidation plugin for this ObjectGrid instance. This plug-in can be used to validate that a javax.security.auth.Subject subject that is passed to the ObjectGrid is a valid subject that has not been tampered with. An implementation of this plug-in needs support from the Subject object creator, because only the creator knows if the Subject object has been tampered with. However, a subject creator might not know if the Subject has been tampered with. In this case, this plug-in should not be used. |
| void setObjectGridAuthorization (ObjectGridAuthorization ogAuthorization); | Sets the ObjectGridAuthorization for this ObjectGrid instance. Passing null to this method removes a previously set ObjectGridAuthorization object from an earlier invocation of this method and indicates that this <code>ObjectGrid</code> is not associated with a ObjectGridAuthorization object. This method should only be used when ObjectGrid security is enabled. If the ObjectGrid security is disabled, the provided ObjectGridAuthorization object will not be used. A ObjectGridAuthorization plugin can be used to authorize access to the ObjectGrid and maps. As of XD 6.1, the setMapAuthorization is deprecated and setObjectGridAuthorization is recommended for use. However, if both MapAuthorization plugin and ObjectGridAuthorization plugin are used, ObjectGrid will use the provided MapAuthorization to authorize map accesses, even though it is deprecated. |

ObjectGrid interface: plug-ins

The ObjectGrid interface has several optional plug-in points for more extensible interactions.

```
void addEventListener(ObjectGridEventListener cb);
void setEventListeners(List cbList);
void removeEventListener(ObjectGridEventListener cb);
void setTransactionCallback(TransactionCallback callback);
int reserveSlot(String);
// Security related plug-ins
void setSubjectValidation(SubjectValidation subjectValidation);
void setSubjectSource(SubjectSource source);
void setMapAuthorization(MapAuthorization mapAuthorization);
```

- **ObjectGridEventListener:** An ObjectGridEventListener interface is used to receive notifications when significant events occur on the ObjectGrid. These events include ObjectGrid initialization, beginning of a transaction, ending a transaction, and destroying an ObjectGrid. To listen for these events, create a class that implements the ObjectGridEventListener interface and add it to the ObjectGrid. These listeners are associated with each Session. See Listeners and Session interface for more information.
- **TransactionCallback:** A TransactionCallback listener interface allows transactional events such as begin, commit and rollback signals to send to this interface. Typically, a TransactionCallback listener interface is used with a Loader. For more information, see TransactionCallback plug-in and Loaders. These events can then be used to coordinate transactions with an external resource or within multiple loaders.
- **reserveSlot:** Allows plug-ins on this ObjectGrid to reserve slots for use in object instances that have slots like TxID.

- **SubjectValidation.** If security is enabled, this plug-in can be used to validate a `javax.security.auth.Subject` class that is passed to the `ObjectGrid`.
- **MapAuthorization.** If security is enabled, this plug-in can be used to authorize `ObjectMap` accesses to the principals that are represented by the `Subject` object.
- **SubjectSource:** If security is enabled, this plug-in can be used to get a `Subject` object that represents the `ObjectGrid` client. This subject is then used for `ObjectGrid` authorization.

For more information about plug-ins, see the introduction to plug-ins in the *Programming Guide*.

BackingMap interface

Each `ObjectGrid` instance contains a collection of `BackingMap` objects. Use the `defineMap` method or the `createMap` method of the `ObjectGrid` interface to name and add each `BackingMap` to an `ObjectGrid` instance. These methods return a `BackingMap` instance that is then used to define the behavior of an individual `Map`. A `BackingMap` can be considered as an in-memory cache of committed data for an individual map.

Session interface

The `Session` interface is used to begin a transaction and to obtain the `ObjectMap` or `JavaMap` that is required for performing transactional interaction between an application and a `BackingMap` object. However, the transaction changes are not applied to the `BackingMap` object until the transaction is committed. A `BackingMap` can be considered as an in-memory cache of committed data for an individual map. For more information, see the information about using `Sessions` to access data in the *Programming Guide*.

The `BackingMap` interface provides methods for setting `BackingMap` attributes. Some of the set methods allow extensibility of a `BackingMap` through several custom designed plug-ins. See the following list of the set methods for setting attributes and providing custom designed plug-in support:

```
// For setting BackingMap attributes.
public void setReadOnly(boolean readOnlyEnabled);
public void setNullValuesSupported(boolean nullValuesSupported);
public void setLockStrategy( LockStrategy lockStrategy );
public void setCopyMode(CopyMode mode, Class valueInterface);
public void setCopyKey(boolean b);
public void setNumberOfBuckets(int numBuckets);
public void setNumberOfLockBuckets(int numBuckets);
public void setLockTimeout(int seconds);
public void setTimeToLive(int seconds);
public void setTtlEvictorType(TTLType type);
public void setEvictionTriggers(String evictionTriggers);

// For setting an optional custom plug-in provided by application.
public abstract void setObjectTransformer(ObjectTransformer t);
public abstract void setOptimisticCallback(OptimisticCallback checker);
public abstract void setLoader(Loader loader);
public abstract void setPreloadMode(boolean async);
public abstract void setEvictor(Evictor e);
public void setMapEventListeners( List /*MapEventListener*/ eventListenerList );
public void addMapEventListener(MapEventListener eventListener );
public void removeMapEventListener(MapEventListener eventListener );
public void addMapIndexPlugin(MapIndexPlugin index);
public void setMapIndexPlugins(List /* MapIndexPlugin */ indexList );
public void createDynamicIndex(String name, boolean isRangeIndex,
```

```
String attributeName, DynamicIndexCallback cb);
public void createDynamicIndex(MapIndexPlugin index, DynamicIndexCallback cb);
public void removeDynamicIndex(String name);
```

A corresponding get method exists for each of the set methods listed.

BackingMap attributes

Each BackingMap has the following attributes that can be set to modify or control the BackingMap behavior:

- **ReadOnly:** This attribute indicates if the Map is a read-only Map or a read and write Map. If this attribute is never set for the Map, then the Map is defaulted to be a read and write Map. When a BackingMap is set to be read only, ObjectGrid optimizes performance for read only when possible.
- **NullValuesSupported:** This attribute indicates if a null value can be put into the Map. If this attribute is never set, the Map does not support null values. If null values are supported by the Map, a get operation that returns null can mean that either the value is null or the map does not contain the key specified by the get operation.
- **LockStrategy:** This attribute determines if a lock manager is used by this BackingMap. If a lock manager is used, then the LockStrategy attribute is used to indicate whether an optimistic locking or pessimistic locking approach is used for locking the map entries. If this attribute is not set, then the optimistic LockStrategy is used. See the Locking topic for details on the supported lock strategies.
- **CopyMode:** This attribute determines if a copy of a value object is made by the BackingMap when a value is read from the map or is put into the BackingMap during the commit cycle of a transaction. Various copy modes are supported to allow the application to make the trade-off between performance and data integrity. If this attribute is not set, then the COPY_ON_READ_AND_COMMIT copy mode is used. This copy mode does not have the best performance, but it has the greatest protection against data integrity problems. If the BackingMap is associated with an EntityManager API entity, then the CopyMode setting has no effect unless the value is set to COPY_TO_BYTES. If any other CopyMode is set, it is always set to NO_COPY. To override the CopyMode for an entity map, use the ObjectMap.setCopyMode method. For more information about the copy modes, see the information about CopyMode method best practices in the *Programming Guide*.
- **CopyKey:** This attribute determines if the BackingMap makes a copy of a key object when an entry is first created in the map. The default action is to not make a copy of key objects because keys are normally unchangeable objects.
- **NumberOfBuckets:** This attribute indicates the number of hash buckets to be used by the BackingMap. The BackingMap implementation uses a hash map for its implementation. If a lot of entries exist in the BackingMap, then more buckets means better performance. The number of keys that have the same bucket becomes lower as the number of buckets grows. More buckets also mean more concurrency. This attribute is useful for fine tuning performance. An undefined default value is used if the application does not set the NumberOfBuckets attribute.
- **NumberOfLockBuckets:** This attribute indicates the number of lock buckets that are be used by the lock manager for this BackingMap. When the LockStrategy is set to OPTIMISTIC or PESSIMISTIC, a lock manager is created for the BackingMap. The lock manager uses a hash map to keep track of entries that are locked by one or more transactions. If a lot of entries exist in the hash map, more lock buckets lead to better performance because the number of keys that

collide on the same bucket is lower as the number of buckets grows. More lock buckets also means more concurrency. When the LockStrategy attribute is set to NONE, no lock manager is used by this BackingMap. In this case, setting numberOfLockBuckets has no effect. If this attribute is not set, an undefined value of is used.

- **LockTimeout:** This attribute is used when the BackingMap is using a lock manager. The BackingMap uses a lock manager when the LockStrategy attribute is set to either OPTIMISTIC or PESSIMISTIC. The attribute value is in seconds and determines how long the lock manager waits for a lock to be granted. If this attribute is not set, then 15 seconds is used as the LockTimeout value. See Pessimistic locking for details regarding the lock wait timeout exceptions that can occur.
- **TtlEvictorType:** Every BackingMap has its own built in time to live evictor that uses a time-based algorithm to determine which map entries to evict. By default, the built in time to live evictor is not active. You can activate the time to live evictor by calling the setTtlEvictorType method with one of these values: CREATION_TIME, LAST_ACCESS_TIME, LAST_UPDATE_TIME, or NONE. A value of CREATION_TIME indicates that the evictor adds the TimeToLive attribute to the time that the map entry was created in the BackingMap to determine when the evictor should evict the map entry from the BackingMap. A value of LAST_ACCESS_TIME (or LAST_UPDATE_TIME) indicates that the evictor adds the TimeToLive attribute to the time that the map entry was last accessed (or accessed *and* updated) by some transaction that the application is running to determine when evictor should evict the map entry. The map entry is evicted only if a map entry is never accessed by any transaction for a period of time that is specified by the TimeToLive attribute. A value of NONE indicates the evictor should remain inactive and never evict any of the map entries. If this attribute is never set, then NONE is used as the default and the time to live evictor is not active. See Evictors for details regarding the built-in time to live evictor.
- **TimeToLive:** This attribute is used to specify the number of seconds that the built in time to live evictor needs to add to the creation or last access time for each entry as described for the TtlEvictorType attribute. If this attribute is never set, then the special value of zero is used to indicate the time to live is infinity. If this attribute is set to infinity, map entries are never evicted by the evictor.

The following example demonstrates how to define the someMap BackingMap in the someGrid ObjectGrid instance and set various attributes of the BackingMap by using the set methods of the BackingMap interface:

```
import com.ibm.websphere.objectgrid.BackingMap;
import com.ibm.websphere.objectgrid.LockStrategy;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;

...

ObjectGrid og =
ObjectGridManagerFactory.getObjectGridManager().createObjectGrid("someGrid");
BackingMap bm = objectGrid.getMap("someMap");
bm.setReadOnly( true ); // override default of read/write
bm.setNullValuesSupported(false); // override default of allowing Null values
bm.setLockStrategy( LockStrategy.PESSIMISTIC ); // override default of OPTIMISTIC
bm.setLockTimeout( 60 ); // override default of 15 seconds.
bm.setNumberOfBuckets(251); // override default (prime numbers work best)
bm.setNumberOfLockBuckets(251); // override default (prime numbers work best)
```


BackingMap plug-ins

The BackingMap interface has several optional plug points for more extensible interactions with the ObjectGrid:

- **ObjectTransformer plug-in:** For some map operations, a BackingMap might need to serialize, deserialize, or copy a key or value of an entry in the BackingMap. The BackingMap can perform these actions by providing a default implementation of the ObjectTransformer interface. An application can improve performance by providing a custom designed ObjectTransformer plug-in that is used by the BackingMap to serialize, deserialize, or copy a key or value of an entry in the BackingMap. See the information about the ObjectTransformer plug-in in the *Programming Guide* for more information.
- **Evictor plug-in:** The built in time to live evictor uses a time-based algorithm to decide when an entry in BackingMap must be evicted. Some applications might need to use a different algorithm for deciding when an entry in a BackingMap needs to be evicted. The Evictor plug-in makes a custom designed Evictor available to the BackingMap to use. The Evictor plug-in is in addition to the built in time to live evictor. It does not replace the time to live evictor. ObjectGrid provides a custom Evictor plug-in that implements well-known algorithms such as "least recently used" or "least frequently used". Applications can either plug-in one of the provided Evictor plug-ins or it can provide its own Evictor plug-in. See the information about eviction in the *Programming Guide*.
- **MapEventListener plug-in:** An application might want to know about BackingMap events such as a map entry eviction or a preload of a BackingMap completion. A BackingMap calls methods on the MapEventListener plug-in to notify an application of BackingMap events. An application can receive notification of various BackingMap events by using the setMapEventListener method to provide one or more custom designed MapEventListener plug-ins to the BackingMap. The application can modify the listed MapEventListener objects by using the addMapEventListener method or the removeMapEventListener method. See the information about the MapEventListener plug-in in the *Programming Guide* for more information.
- **Loader plug-in:** A BackingMap is an in-memory cache of a Map. A Loader plug-in is an option that is used by the BackingMap to move data between memory and a persistent store.. For example, a Java database connectivity (JDBC) Loader can be used to move data in and out of a BackingMap and one or more relational tables of a relational database. A relational database does not need to be used as the persistent store for a BackingMap. The Loader can also be used to moved data between a BackingMap and a file, between a BackingMap and a Hibernate map, between a BackingMap and a Java 2 Platform, Enterprise Edition (JEE) entity bean, between a BackingMap and another application server, and so on. The application must provide a custom-designed Loader plug-in to move data between the BackingMap and the persistent store for every technology that is used. If a Loader is not provided, the BackingMap becomes a simple in-memory cache. See the information about using a Loader in the *Programming Guide* for more information.
- **OptimisticCallback plug-in:** When the LockStrategy attribute for a BackingMap is set to OPTIMISTIC, either the BackingMap or a Loader plug-in must perform comparison operations for the values of the map. The OptimisticCallback plug-in is used by the BackingMap and the Loader to perform the optimistic versioning comparison operations. See the information about the OptimisticCallback plug-in in the *Programming Guide* for more information.
- **MapIndexPlugin plug-in:** A MapIndexPlugin plug-in, or an Index in short, is an option that is used by the BackingMap to build an index that is based on the

specified attribute of the stored object. The index allows the application to find objects by a specific value or a range of values. There are two types of index: static and dynamic. Refer to Indexing for detailed information.

For more information regarding plug-ins, see the introduction to plug-ins in the *Programming Guide*.

Connecting to a distributed ObjectGrid

You can connect to a distributed ObjectGrid with a connection end point for the catalog service domain. You must have the host name and listener port of each catalog server in the catalog service domain to which you want to connect.

To connect to a distributed data grid, you must have configured your server-side environment with a catalog service and container servers.

For more information about finding the listener port for each catalog service, refer to .Read more about finding the listener port for each catalog service in the *Administration Guide*.

The `getObjectGrid(ClientClusterContext ccc, String objectGridName)` method connects to the specified catalog service domain and returns a client ObjectGrid instance corresponding to a server-side ObjectGrid instance.

The following code allows you to connect to a stand-alone distributed data grid using explicit catalog service end points:

```
// Create an ObjectGridManager instance.

ObjectGridManager ogm = ObjectGridManagerFactory.getObjectGridManager();

// Obtain a ClientClusterContext by connecting to a catalog
// server based distributed ObjectGrid. You have to provide
// a connection end point for your catalog server in the format
// of hostName:endPointPort. The hostName is the machine
// where the catalog server resides, and the endPointPort is
// the catalog server's listening port, whose default is 2809.
// Catalog server end-points for a given domain must be in
// the format of a comma-delimited list.

String catalogServiceEndpoints = "host1:2809,host2:2809";
ClientClusterContext ccc = ogm.connect(catalogServiceEndpoints, null, null);

// Obtain a distributed ObjectGrid using ObjectGridManager and providing
// the ClientClusterContext.

ObjectGrid og = ogm.getObjectGrid(ccc, "objectdata gridName");
```

To connect to a catalog service domain from a client application hosted in WebSphere Application Server, where the catalog service domain has been configured using the administrative console or admin task, the catalog service endpoints can be retrieved using the embedded server API:

```
...

// Retrieve the catalog service endpoints from the ServerProperties
// singleton, which is configured in the WebSphere administration
// console or admin task.

String catalogServiceEndpoints = ServerFactory.getServerProperties()
    .getCatalogServiceBootstrap();
```

```
ClientClusterContext ccc = ogm.connect(catalogServiceEndpoints,  
    null, null);  
  
...
```

If the catalog service domain in WebSphere Application Server is hosted by the deployment manager, clients outside of the cell, including Java Platform, Standard Edition clients, must connect to the catalog service using the deployment manager host name and the IIOP bootstrap port. When the catalog service runs in WebSphere Application Server cells while the clients run outside of the cells, look to the eXtreme Scale domain configuration pages in the WebSphere Application Server administrative console for the information needed to point a client to the catalog service.

Interacting with an ObjectGrid using ObjectGridManager

The ObjectGridManagerFactory class and the ObjectGridManager interface provide a mechanism to create, access, and cache ObjectGrid instances. The ObjectGridManagerFactory class is a static helper class to access the ObjectGridManager interface, a singleton. The ObjectGridManager interface includes several convenience methods to create instances of an ObjectGrid object. The ObjectGridManager interface also facilitates creation and caching of ObjectGrid instances that can be accessed by several users.

Programming model

Before using eXtreme Scale's functionality as an in-memory data grid, you must create and interact with ObjectGrid instances with methods such as the following.

- createObjectGrid methods
- getObjectGrid methods
- removeObjectGrid methods
- controlling the life cycle of an ObjectGrid

createObjectGrid methods

This topic describes the seven createObjectGrid methods in the ObjectGridManager interface. Each of these methods creates a local instance of an ObjectGrid.

Local in-memory instance

The following code snippet illustrates how to obtain and configure a local ObjectGrid instance with eXtreme Scale.

```
// Obtain a local ObjectGrid reference  
// you can create a new ObjectGrid, or get configured ObjectGrid  
// defined in ObjectGrid xml file  
ObjectGridManager objectGridManager =  
ObjectGridManagerFactory.getObjectGridManager();  
ObjectGrid ivObjectGrid =  
objectGridManager.createObjectGrid("objectgridName");  
  
// Add a TransactionCallback into ObjectGrid  
HeapTransactionCallback tcb = new HeapTransactionCallback();  
ivObjectGrid.setTransactionCallback(tcb);  
  
// Define a BackingMap  
// if the BackingMap is configured in ObjectGrid xml  
// file, you can just get it.  
BackingMap ivBackingMap = ivObjectGrid.defineMap("myMap");
```

```

// Add a Loader into BackingMap
Loader ivLoader = new HeapCacheLoader();
ivBackingMap.setLoader(ivLoader);

// initialize ObjectGrid
ivObjectGrid.initialize();

// Obtain a session to be used by the current thread.
// Session can not be shared by multiple threads
Session ivSession = ivObjectGrid.getSession();

// Obtaining ObjectMap from ObjectGrid Session
ObjectMap objectMap = ivSession.getMap("myMap");

```

Default shared configuration

The following code is a simple case of creating an ObjectGrid to share among many users.

```

import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.ObjectGridException;
import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
import com.ibm.websphere.objectgrid.ObjectGridManager;
final ObjectGridManager oGridManager=
    ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid employees =
    oGridManager.createObjectGrid("Employees",true);
employees.initialize();
employees.
/*sample continues..*/

```

The preceding Java code snippet creates and caches the Employees ObjectGrid. The Employees ObjectGrid is initialized with the default configuration and is ready for use. The second parameter in the createObjectGrid method is set to true, which instructs the ObjectGridManager to cache the ObjectGrid instance it creates. If this parameter is set to false, the instance is not cached. Every ObjectGrid instance has a name, and the instance can be shared among many clients or users based on that name.

If the objectGrid instance is used in peer-to-peer sharing, the caching must be set to true. For more information on peer-to-peer sharing, see [Distributing changes between peer Java Virtual Machines](#).

XML configuration

WebSphere eXtreme Scale is highly configurable. The previous example demonstrates how to create a simple ObjectGrid without any configuration. This example shows you how to create a pre-configured ObjectGrid instance that is based on an XML configuration file. You can configure an ObjectGrid instance programmatically or using an XML-based configuration file. You can also configure ObjectGrid using a combination of both approaches. The ObjectGridManager interface allows creation of an ObjectGrid instance based on the XML configuration. The ObjectGridManager interface has several methods that take a URL as an argument. Every XML file that is passed into the ObjectGridManager must be validated against the schema. XML validation can be disabled only when the file is previously validated and no changes have been made to the file since its last validation. Disabling validation saves a small amount of overhead but introduces the possibility of using an invalid XML file. The IBM Java Developer Kit (JDK) 1.4.2 has support for XML validation. When using a JDK that does not have this support, Apache Xerces might be required to validate the XML.

The following Java code snippet demonstrates how to pass in an XML configuration file to create an ObjectGrid.

```
import java.net.MalformedURLException;
import java.net.URL;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.ObjectGridException;
import com.ibm.websphere.objectgrid.ObjectGridManager;
import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
boolean validateXML = true; // turn XML validation on
boolean cacheInstance = true; // Cache the instance
String objectGridName="Employees"; // Name of Object Grid URL
allObjectGrids = new URL("file:test/myObjectGrid.xml");
final ObjectGridManager oGridManager=
    ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid employees =
    oGridManager.createObjectGrid(objectGridName, allObjectGrids,
        bvalidateXML, cacheInstance);
```

The XML file can contain configuration information for several ObjectGrids. The previous code snippet specifically returns ObjectGrid Employees, assuming that the Employees configuration is defined in the file. For the XML syntax, see ObjectGrid configuration. Seven createObjectGrid methods exist, and are documented in the following code block.

```
/**
 * A simple factory method to return an instance of an
 * Object Grid. A unique name is assigned.
 * The instance of ObjectGrid is not cached.
 * Users can then use {@link ObjectGrid#setName(String)} to change the
 * ObjectGrid name.
 *
 * @return ObjectGrid an instance of ObjectGrid with a unique name assigned
 * @throws ObjectGridException any error encountered during the
 * ObjectGrid creation
 */
public ObjectGrid createObjectGrid() throws ObjectGridException;

/**
 * A simple factory method to return an instance of an ObjectGrid with the
 * specified name. The instances of ObjectGrid can be cached. If an ObjectGrid
 * with the this name has already been cached, an ObjectGridException
 * will be thrown.
 *
 * @param objectGridName the name of the ObjectGrid to be created.
 * @param cacheInstance true, if the ObjectGrid instance should be cached
 * @return an ObjectGrid instance
 * @this name has already been cached or
 * any error during the ObjectGrid creation.
 */
public ObjectGrid createObjectGrid(String objectGridName, boolean cacheInstance)
    throws ObjectGridException;

/**
 * Create an ObjectGrid instance with the specified ObjectGrid name. The
 * ObjectGrid instance created will be cached.
 * @param objectGridName the Name of the ObjectGrid instance to be created.
 * @return an ObjectGrid instance
 * @throws ObjectGridException if an ObjectGrid with this name has already
 * been cached, or any error encountered during the ObjectGrid creation
 */
public ObjectGrid createObjectGrid(String objectGridName)
    throws ObjectGridException;

/**
 * Create an ObjectGrid instance based on the specified ObjectGrid name and the
 * XML file. The ObjectGrid instance defined in the XML file with the specified
```

```

* ObjectGrid name will be created and returned. If such an ObjectGrid
* cannot be found in the xml file, an exception will be thrown.
*
* This ObjectGrid instance can be cached.
*
* If the URL is null, it will be simply ignored. In this case, this method behaves
* the same as {@link #createObjectGrid(String, boolean)}.
*
* @param objectGridName the Name of the ObjectGrid instance to be returned. It
* must not be null.
* @param xmlFile a URL to a wellformed xml file based on the ObjectGrid schema.
* @param enableXmlValidation if true the XML is validated
* @param cacheInstance a boolean value indicating whether the ObjectGrid
* instance(s)
* defined in the XML will be cached or not. If true, the instance(s) will
* be cached.
*
* @throws ObjectGridException if an ObjectGrid with the same name
* has been previously cached, no ObjectGrid name can be found in the xml file,
* or any other error during the ObjectGrid creation.
* @return an ObjectGrid instance
* @see ObjectGrid
*/
public ObjectGrid createObjectGrid(String objectGridName, final URL xmlFile,
final boolean enableXmlValidation, boolean cacheInstance)
throws ObjectGridException;

/**
* Process an XML file and create a List of ObjectGrid objects based
* upon the file.
* These ObjectGrid instances can be cached.
* An ObjectGridException will be thrown when attempting to cache a
* newly created ObjectGrid
* that has the same name as an ObjectGrid that has already been cached.
*
* @param xmlFile the file that defines an ObjectGrid or multiple
* ObjectGrids
* @param enableXmlValidation setting to true will validate the XML
* file against the schema
* @param cacheInstances set to true to cache all ObjectGrid instances
* created based on the file
* @return an ObjectGrid instance
* @throws ObjectGridException if attempting to create and cache an
* ObjectGrid with the same name as
* an ObjectGrid that has already been cached, or any other error
* occurred during the
* ObjectGrid creation
*/
public List createObjectGrids(final URL xmlFile, final boolean enableXmlValidation,
boolean cacheInstances) throws ObjectGridException;

/** Create all ObjectGrids that are found in the XML file. The XML file will be
* validated against the schema. Each ObjectGrid instance that is created will
* be cached. An ObjectGridException will be thrown when attempting to cache a
* newly created ObjectGrid that has the same name as an ObjectGrid that has
* already been cached.
* @param xmlFile The XML file to process. ObjectGrids will be created based
* on what is in the file.
* @return A List of ObjectGrid instances that have been created.
* @throws ObjectGridException if an ObjectGrid with the same name as any of
* those found in the XML has already been cached, or
* any other error encountered during ObjectGrid creation.
*/
public List createObjectGrids(final URL xmlFile) throws ObjectGridException;

/**
* Process the XML file and create a single ObjectGrid instance with the

```

```

* objectGridName specified only if an ObjectGrid with that name is found in
* the file. If there is no ObjectGrid with this name defined in the XML file,
* an ObjectGridException
* will be thrown. The ObjectGrid instance created will be cached.
* @param objectGridName name of the ObjectGrid to create. This ObjectGrid
* should be defined in the XML file.
* @param xmlFile the XML file to process
* @return A newly created ObjectGrid
* @throws ObjectGridException if an ObjectGrid with the same name has been
* previously cached, no ObjectGrid name can be found in the xml file,
* or any other error during the ObjectGrid creation.
*/
public ObjectGrid createObjectGrid(String objectGridName, URL xmlFile)
    throws ObjectGridException;

```

Client hangs during a getObjectGrid method call

A client might seem to hang when calling the getObjectGrid method on the ObjectGridManager or throw an exception: com.ibm.websphere.projector.MetadataException. The EntityMetadata repository is not available and the timeout threshold is reached. The reason is the client is waiting for the entity metadata on the ObjectGrid server to become available. This error can occur when a container has been started, but the initial number of containers or minimum number of synchronous replicas has not been reached. Examine the deployment policy for the ObjectGrid and verify that the number of active containers is greater than or equal to both the numInitialContainers and minSyncReplicas attributes in the deployment policy descriptor file.

getObjectGrid methods

Use the ObjectGridManager.getObjectGrid methods to retrieve cached instances.

Retrieving a cached instance

Since the Employees ObjectGrid instance was cached by the ObjectGridManager interface, another user can access it with the following code snippet:

```
ObjectGrid myEmployees = oGridManager.getObjectGrid("Employees");
```

The following are the two getObjectGrid methods that return cached ObjectGrid instances:

- **Retrieving all cached instances**

To obtain all of the ObjectGrid instances that have been previously cached, use the getObjectGrids method, which returns a list of each instance. If no cached instances exist, the method will return null.

- **Retrieving a cached instance by name**

To obtain a single cached instance of an ObjectGrid, use getObjectGrid(String objectGridName), passing the name of the cached instance into the method. The method either returns the ObjectGrid instance with the specified name or returns null if there is no ObjectGrid instance with that name.

Note: You can also use the getObjectGrid method to connect to a distributed grid. See “Connecting to a distributed ObjectGrid” on page 16 for more information.

removeObjectGrid methods

You can use two different removeObjectGrid methods to remove ObjectGrid instances from the cache.

Remove an ObjectGrid instance

To remove ObjectGrid instances from the cache, use one of the removeObjectGrid methods. The ObjectGridManager interface does not keep a reference of the instances that are removed. Two remove methods exist. One method takes a boolean parameter. If the boolean parameter is set to true, the destroy method is called on the ObjectGrid. The call to the destroy method on the ObjectGrid shuts down the ObjectGrid and frees up any resources the ObjectGrid is using. A description of how to use the two removeObjectGrid methods follows:

```
/**
 * Remove an ObjectGrid from the cache of ObjectGrid instances
 *
 * @param objectGridName the name of the ObjectGrid instance to remove
 * from the cache
 *
 * @throws ObjectGridException if an ObjectGrid with the objectGridName
 * was not found in the cache
 */
public void removeObjectGrid(String objectGridName) throws ObjectGridException;

/**
 * Remove an ObjectGrid from the cache of ObjectGrid instances and
 * destroy its associated resources
 *
 * @param objectGridName the name of the ObjectGrid instance to remove
 * from the cache
 *
 * @param destroy destroy the objectgrid instance and its associated
 * resources
 *
 * @throws ObjectGridException if an ObjectGrid with the objectGridName
 * was not found in the cache
 */
public void removeObjectGrid(String objectGridName, boolean destroy)
    throws ObjectGridException;
```

Controlling the life cycle of an ObjectGrid

You can use the ObjectGridManager interface to control the life cycle of an ObjectGrid instance using either a startup bean or a servlet.

Managing life cycle with a startup bean

A startup bean is used to control the life cycle of an ObjectGrid instance. A startup bean loads when an application starts. With a startup bean, code can run whenever an application starts or stops as expected. To create a startup bean, use the home.com.ibm.websphere.startupservice.AppStartUpHome interface and use the remote.com.ibm.websphere.startupservice.AppStartUp interface. Implement the start and stop methods on the bean. The start method is invoked whenever the application starts up. The stop method is invoked when the application shuts down. The start method is used to create ObjectGrid instances. The stop method is used to remove ObjectGrid instances. A code snippet that demonstrates this ObjectGrid life-cycle management in a startup bean follows:

```
public class MyStartupBean implements javax.ejb.SessionBean {
    private ObjectGridManager objectGridManager;

    /* The methods on the SessionBean interface have been
     * left out of this example for the sake of brevity */

    public boolean start(){
        // Starting the startup bean
        // This method is called when the application starts
```



```

        objectGridManager = ObjectGridManagerFactory.getObjectGridManager();
        try {
            // create 2 ObjectGrids and cache these instances
            ObjectGrid bookstoreGrid =
objectGridManager.createObjectGrid("bookstore", true);
            bookstoreGrid.defineMap("book");
            ObjectGrid videostoreGrid =
objectGridManager.createObjectGrid("videostore", true);
            // within the JVM,
            // these ObjectGrids can now be retrieved from the
            //ObjectGridManager using the getObjectGrid(String) method
        } catch (ObjectGridException e) {
            e.printStackTrace();
            return false;
        }

        return true;
    }

    public void stop(){
        // Stopping the startup bean
        // This method is called when the application is stopped
        try {
            // remove the cached ObjectGrids and destroy them
            objectGridManager.removeObjectGrid("bookstore", true);
            objectGridManager.removeObjectGrid("videostore", true);
        } catch (ObjectGridException e) {
            e.printStackTrace();
        }
    }
}

```

After the start method is called, the newly created ObjectGrid instances are retrieved from the ObjectGridManager interface. For example, if a servlet is included in the application, the servlet accesses the eXtreme Scale using the following code snippet:

```

ObjectGridManager objectGridManager =
    ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid bookstoreGrid = objectGridManager.getObjectGrid("bookstore");
ObjectGrid videostoreGrid = objectGridManager.getObjectGrid("videostore");

```

Managing life cycle with a servlet

To manage the life cycle of an ObjectGrid in a servlet, you can use the init method to create an ObjectGrid instance and the destroy method to remove the ObjectGrid instance. If the ObjectGrid instance is cached, it is retrieved and manipulated in the servlet code. Sample code that demonstrates ObjectGrid creation, manipulation, and destruction within a servlet follows:

```

public class MyObjectGridServlet extends HttpServlet implements Servlet {
    private ObjectGridManager objectGridManager;

    public MyObjectGridServlet() {
        super();
    }

    public void init(ServletConfig arg0) throws ServletException {
        super.init();
        objectGridManager = ObjectGridManagerFactory.getObjectGridManager();
        try {
            // create and cache an ObjectGrid named bookstore
            ObjectGrid bookstoreGrid =
objectGridManager.createObjectGrid("bookstore", true);
            bookstoreGrid.defineMap("book");
        } catch (ObjectGridException e) {

```

```

        e.printStackTrace();
    }
}

protected void doGet(HttpServletRequest req, HttpServletResponse res)
    throws ServletException, IOException {
    ObjectGrid bookstoreGrid = objectGridManager.getObjectGrid("bookstore");
    Session session = bookstoreGrid.getSession();
    ObjectMap bookMap = session.getMap("book");
    // perform operations on the cached ObjectGrid
    // ...
}

public void destroy() {
    super.destroy();
    try {
        // remove and destroy the cached bookstore ObjectGrid
        objectGridManager.removeObjectGrid("bookstore", true);
    } catch (ObjectGridException e) {
        e.printStackTrace();
    }
}
}
}

```

Accessing the ObjectGrid shard

WebSphere eXtreme Scale achieves high processing rates by moving the logic to where the data is and returning only results back to the client.

Application logic in a client Java virtual machine (JVM) needs to pull data from the server JVM that is holding the data and push it back when the transaction commits. This process slows down the rate the data can be processed. If the application logic was on the same JVM as the shard that is holding the data, then the network latency and marshalling cost is eliminated and can provide a significant performance boost.

Local reference to shard data

The ObjectGrid APIs provide a Session to the server-side method. This session is a direct reference to the data for that shard. No routing logic is on that path. The application logic can work with the data for that shard directly. The session cannot be used to access data in another partition because no routing logic exists.

A Loader plug-in also provides a way to receive an event when a shard becomes a primary partition. An application can implement a Loader and implement the ReplicaPreloadController interface. The check preload status method is only called when a shard becomes a primary. The session provided to that method is a local reference to the shards data. This approach is typically used if a partition primary needs to start some threads or subscribe to a message fabric for partition-related traffic. It might start a thread to listen for messages in a local Map using the getNextKey API.

Collocated client-server optimization

If an application uses the client APIs to access a partition that happens to be collocated with the JVM that contains the client, then the network is avoided but some marshalling still occurs because of current implementation issues. If a partitioned grid is used, then no impact on the performance of the application is made because (N-1)/N number of calls route to a different JVM. If you need local

access always with a shard, then use the Loader or ObjectGrid APIs to invoke that logic.

Data access with indexes (Index API)

Using indexing as an alternative to key access for data can be more efficient.

Necessary steps

1. Add either static or dynamic index plug-ins to the BackingMap.
2. Obtain application index proxy object by issuing the `getIndex` method of `ObjectMap`.
3. Cast the index proxy object to an appropriate application index interface, such as `MapIndex`, `MapRangeIndex`, or a customized index interface.
4. Use the methods defined in the application index interface to find cached objects.

The `HashIndex` class is the built-in index plug-in implementation that can support both of the built-in application index interfaces: `MapIndex` and `MapRangeIndex`. You also can create your own indexes. You can add `HashIndex` as either a static or dynamic index into the `BackingMap`, obtain either `MapIndex` or `MapRangeIndex` index proxy object, and use the index proxy object to find cached objects.

Note: In a distributed environment, if the index object is obtained from a client `ObjectGrid`, it will have type client index object and all index operations will run in a remote server `ObjectGrid`. If the map is partitioned, the index operation will run on each partition remotely and results from each partition will be merged before returning them to the application. The performance will be determined by the number of partitions and the size of the result returned by each partition. Poor performance may occur if both factors are high.

For information about configuring the `HashIndex`, refer to [Configuring the HashIndex](#).

If you want to write your own index plug-in, see [“Plug-ins for custom indexing of cache objects”](#) on page 30.

For information regarding indexing, see [Indexing and “Using a composite index”](#) on page 227.

Adding static index plug-ins

You can use two approaches to add static index plug-ins into the `BackingMap` configuration: XML configuration and programmatic configuration. The following example illustrates the XML configuration approach.

Adding static index plug-ins: XML configuration approach

```
<backingMapPluginCollection id="person">
  <bean id="MapIndexplugin"
    className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
    <property name="Name" type="java.lang.String" value="CODE"
      description="index name" />
    <property name="RangeIndex" type="boolean" value="true"
      description="true for MapRangeIndex" />
    <property name="AttributeName" type="java.lang.String" value="employeeCode"
      description="attribute name" />
  </bean>
</backingMapPluginCollection>
```

In this XML configuration example, the built-in HashIndex class is used as the index plug-in. The HashIndex supports properties that users can configure, such as Name, RangeIndex, and AttributeName in the previous example.

- The Name property is configured as “CODE”, a string identifying this index plug-in. The Name property value must be unique within the scope of the BackingMap, and can be used to retrieve the index object by name from the ObjectMap instance for the BackingMap.
- The RangeIndex property is configured as “true”, which means the application can cast the retrieved index object to the MapRangeIndex interface. If the RangeIndex property is configured as “false”, the application can only cast the retrieved index object to the MapIndex interface. A MapRangeIndex supports functions to find data using range functions such as greater than, less than, or both, while a MapIndex only supports equals functions. If the index will be used by query, the RangeIndex property must be configured to “true” on single-attribute indexes. For a relationship index and composite index, the RangeIndex property must be configured to “false”.
- The AttributeName property is configured as “employeeCode”, which means the employeeCode attribute of the cached object is used to build a single-attribute index. If an application needs to search for cached objects with multiple attributes, the AttributeName property can be set to a comma-delimited list of attributes, yielding a composite index.

See the information about configuring HashIndex in the *Administration Guide* for more information.

The BackingMap interface has two methods that you can use to add static index plug-ins: addMapIndexplugin and setMapIndexplugins. For more information, see the API documentation.

The following code example illustrates the programmatic configuration approach:

Adding static index plugins: programmatic configuration approach

```
import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
import com.ibm.websphere.objectgrid.ObjectGridManager;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.BackingMap;

ObjectGridManager ogManager = ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid ivObjectGrid = ogManager.createObjectGrid( "grid" );
BackingMap personBackingMap = ivObjectGrid.getMap("person");

// use the builtin HashIndex class as the index plugin class.
HashIndex mapIndexplugin = new HashIndex();
mapIndexplugin.setName("CODE");
mapIndexplugin.setAttributeName("EmployeeCode");
mapIndexplugin.setRangeIndex(true);
personBackingMap.addMapIndexplugin(mapIndexplugin);
```

Using static indexes

After a static index plug-in is added to a BackingMap configuration and the containing ObjectGrid instance is initialized, applications can retrieve the index object by name from the ObjectMap instance for the BackingMap. Cast the index object to the application index interface. Operations that the application index interface supports can now run.

The following code example illustrates how to retrieve and use static indexes.

Using static indexes example

```

Session session = ivObjectGrid.getSession();
ObjectMap map = session.getMap("person ");
MapRangeIndex codeIndex = (MapRangeIndex) m.getIndex("CODE");
Iterator iter = codeIndex.findLessEqual(new Integer(15));
while (iter.hasNext()) {
    Object key = iter.next();
    Object value = map.get(key);
}

```

Adding, removing, and using dynamic indexes

You can create and remove dynamic indexes from a BackingMap instance programmatically at any time. A dynamic index differs from a static index in that the dynamic index can be created even after the containing ObjectGrid instance is initialized. Unlike static indexing, the dynamic indexing is an asynchronous process and needs to be in ready state before you use it. This method uses the same approach for retrieving and using the dynamic indexes as static indexes. You can remove a dynamic index if it is no longer needed. The BackingMap interface has methods to create and remove dynamic indexes.

See the BackingMap API for more information about the createDynamicIndex and removeDynamicIndex methods.

Using dynamic indexes example

```

import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
import com.ibm.websphere.objectgrid.ObjectGridManager;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.BackingMap;

ObjectGridManager ogManager = ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid og = ogManager.createObjectGrid("grid");
BackingMap bm = og.getMap("person");
og.initialize();

// create index after ObjectGrid initialization without DynamicIndexCallback.
bm.createDynamicIndex("CODE", true, "employeeCode", null);

try {
    // If not using DynamicIndexCallback, need to wait for the Index to be ready.
    // The waiting time depends on the current size of the map
    Thread.sleep(3000);
} catch (Throwable t) {
    // ...
}

// When the index is ready, applications can try to get application index
// interface instance.
// Applications have to find a way to ensure that the index is ready to use,
// if not using DynamicIndexCallback interface.
// The following example demonstrates the way to wait for the index to be ready
// Consider the size of the map in the total waiting time.

Session session = og.getSession();
ObjectMap m = session.getMap("person");
MapRangeIndex codeIndex = null;

int counter = 0;
int maxCounter = 10;
boolean ready = false;
while (!ready && counter < maxCounter) {
    try {
        counter++;
        codeIndex = (MapRangeIndex) m.getIndex("CODE");
        ready = true;
    } catch (IndexNotReadyException e) {
        // implies index is not ready, ...
        System.out.println("Index is not ready. continue to wait.");
        try {
            Thread.sleep(3000);
        } catch (Throwable tt) {
            // ...
        }
    } catch (Throwable t) {
        // unexpected exception
        t.printStackTrace();
    }
}

if (!ready) {
    System.out.println("Index is not ready. Need to handle this situation.");
}

// Use the index to perform queries
// Refer to the MapIndex or MapRangeIndex interface for supported operations.
// The object attribute on which the index is created is the EmployeeCode.
// Assume that the EmployeeCode attribute is Integer type: the
// parameter that is passed into index operations has this data type.

```

```

Iterator iter = codeIndex.findLessEqual(new Integer(15));
// remove the dynamic index when no longer needed
bm.removeDynamicIndex("CODE");

```

DynamicIndexCallback interface

The DynamicIndexCallback interface is designed for applications that want to get notifications at the indexing events of ready, error, or destroy. The DynamicIndexCallback is an optional parameter for the createDynamicIndex method of the BackingMap. With a registered DynamicIndexCallback instance, applications can run business logic upon receiving notification of an indexing event. For example, the ready event means that the index is ready for use. When a notification for this event is received, an application can try to retrieve and use the application index interface instance. See the DynamicIndexCallback API in the API documentation for more information.

The following code example illustrates the use of the DynamicIndexCallback interface:

Using DynamicIndexCallback interface

```

BackingMap personBackingMap = ivObjectGrid.getMap("person");
DynamicIndexCallback callback = new DynamicIndexCallbackImpl();
personBackingMap.createDynamicIndex("CODE", true, "employeeCode", callback);

class DynamicIndexCallbackImpl implements DynamicIndexCallback {
    public DynamicIndexCallbackImpl() {
    }

    public void ready(String indexName) {
        System.out.println("DynamicIndexCallbackImpl.ready() -> indexName = " + indexName);

        // Simulate what an application would do when notified that the index is ready.
        // Normally, the application would wait until the ready state is reached and then proceed
        // with any index usage logic.
        if("CODE".equals(indexName)) {
            ObjectGridManager ogManager = ObjectGridManagerFactory.getObjectGridManager();
            ObjectGrid og = ogManager.createObjectGrid("grid");
            Session session = og.getSession();
            ObjectMap map = session.getMap("person");
            MapIndex codeIndex = (MapIndex) map.getIndex("CODE");
            Iterator iter = codeIndex.findAll(codeValue);
        }
    }

    public void error(String indexName, Throwable t) {
        System.out.println("DynamicIndexCallbackImpl.error() -> indexName = " + indexName);
        t.printStackTrace();
    }

    public void destroy(String indexName) {
        System.out.println("DynamicIndexCallbackImpl.destroy() -> indexName = " + indexName);
    }
}

```

Indexing

Use the MapIndexPlugin to build an index or several indexes on a BackingMap to support non-key data access.

Index types and configuration

The indexing feature is represented by the MapIndexPlugin or Index for short. The Index is a BackingMap plug-in. A BackingMap can have multiple Index plug-ins configured, as long as each one follows the Index configuration rules.

You can use the indexing feature to build an index or several indexes on a BackingMap. An index is built from an attribute or a list of attributes of an object in the BackingMap. This feature provides a way for applications to find certain objects more quickly. With the indexing feature, applications can find objects with a specific value or within a range of values of indexed attributes.

Two types of indexing are possible: static and dynamic. With static indexing, you must configure the index plug-in on the BackingMap before initializing the ObjectGrid instance. You can do this configuration with XML or programmatic configuration of the BackingMap. Static indexing starts building an index during ObjectGrid initialization. The index is always synchronized with the BackingMap and ready for use. After the static indexing process starts, the maintenance of the index is part of the eXtreme Scale transaction management process. When transactions commit changes, these changes also update the static index, and index changes are rolled back if the transaction is rolled back.

With dynamic indexing, you can create an index on a BackingMap before or after the initialization of the containing ObjectGrid instance. Applications have life cycle control over the dynamic indexing process so that you can remove a dynamic index when it is no longer needed. When an application creates a dynamic index, the index might not be ready for immediate use because of the time it takes to complete the index building process. Because the amount of time depends upon the amount of data indexed, the DynamicIndexCallback interface is provided for applications that want to receive notifications when certain indexing events occur. These events include ready, error, and destroy. Applications can implement this callback interface and register with the dynamic indexing process.

If a BackingMap has an index plug-in configured, you can obtain the application index proxy object from the corresponding ObjectMap. Calling the getIndex method on the ObjectMap and passing in the name of the index plug-in returns the index proxy object. You must cast the index proxy object to an appropriate application index interface, such as MapIndex, MapRangeIndex, or a customized index interface. After obtaining the index proxy object, you can use methods defined in the application index interface to find cached objects.

The steps to use indexing are summarized in the following list:

- Add either static or dynamic index plug-ins into the BackingMap.
- Obtain an application index proxy object by issuing the getIndex method of the ObjectMap.
- Cast the index proxy object to an appropriate application index interface, such as MapIndex, MapRangeIndex, or a customized index interface.
- Use methods that are defined in application index interface to find cached objects.

The HashIndex class is the built-in index plug-in implementation that can support both of the built-in application index interfaces: MapIndex and MapRangeIndex. You also can create your own indexes. You can add HashIndex as either a static or dynamic index into the BackingMap, obtain either MapIndex or MapRangeIndex index proxy object, and use the index proxy object to find cached objects.

For information about configuring the HashIndex, refer to *Configuring the HashIndex*.

For more information about writing your own index plug-in, see the information about writing an index plug-in in the *Programming Guide*.

For information about how to use indexing, see the information about using indexing for non-key data access in the *Programming Guide* and *Composite HashIndex*.

Data quality consideration

The results of index query methods only represent a snapshot of data at a point of time. No locks against data entries are obtained after the results return to the application. Application has to be aware that data updates may occur on a returned data set. For example, the application obtains the key of a cached object by running the `findAll` method of `MapIndex`. This returned key object is associated with a data entry in the cache. The application should be able to run the `get` method on `ObjectMap` to find an object by providing the key object. If another transaction removes the data object from the cache just before the `get` method is called, the returned result will be null.

Indexing performance considerations

One of the main objectives of the indexing feature is to improve overall `BackingMap` performance. If indexing is not used properly, the performance of the application might be compromised. Consider the following factors before using this feature.

- **The number of concurrent write transactions:** Index processing can occur every time a transaction writes data into a `BackingMap`. Performance degrades if many transactions are writing data into the map concurrently when an application attempts index query operations.
- **The size of the result set that is returned by a query operation:** As the size of the resultset increases, the query performance declines. Performance tends to degrade when the size of the result set is 15% or more of the `BackingMap`.
- **The number of indexes built over the same `BackingMap`:** Each index consumes system resources. As the number of the indexes built over the `BackingMap` increases, performance decreases.

The indexing function can improve `BackingMap` performance drastically. Ideal cases are when the `BackingMap` has mostly read operations, the query result set is of a small percentage of the `BackingMap` entries, and only few indexes are built over the `BackingMap`.

Plug-ins for custom indexing of cache objects

With a `MapIndexPlugin` plug-in, or `index`, you can write custom indexing strategies that are beyond the built-in indexes that eXtreme Scale provides.

For general information about indexing, see [Indexing](#).

For information about using indexing, see “Data access with indexes (Index API)” on page 25.

`MapIndexPlugin` implementations must use the `MapIndexPlugin` interface and follow the common eXtreme Scale plug-in conventions.

The following sections include some of the important methods of the `index` interface.

setProperties method

Use the `setProperties` method to initialize the index plug-in programmatically. The `Properties` object parameter that is passed into the method should contain required configuration information to initialize the index plug-in properly. The `setProperties` method implementation, along with the `getProperties` method implementation, are

required in a distributed environment because the index plug-in configuration moves between client and server processes. An implementation example of this method follows.

```

setProperties(Properties properties)

// setProperties method sample code
public void setProperties(Properties properties) {
    ivIndexProperties = properties;

    String ivRangeIndexString = properties.getProperty("rangeIndex");
    if (ivRangeIndexString != null && ivRangeIndexString.equals("true")) {
        setRangeIndex(true);
    }
    setName(properties.getProperty("indexName"));
    setAttributeName(properties.getProperty("attributeName"));

    String ivFieldAccessAttributeString = properties.getProperty("fieldAccessAttribute");
    if (ivFieldAccessAttributeString != null && ivFieldAccessAttributeString.equals("true")) {
        setFieldAccessAttribute(true);
    }

    String ivPOJOKeyIndexString = properties.getProperty("POJOKeyIndex");
    if (ivPOJOKeyIndexString != null && ivPOJOKeyIndexString.equals("true")) {
        setPOJOKeyIndex(true);
    }
}

```

getProperties method

The `getProperties` method extracts the index plug-in configuration from a `MapIndexPlugin` instance. You can use the extracted properties to initialize another `MapIndexPlugin` instance to have the same internal states. The `getProperties` method and `setProperties` method implementations are required in a distributed environment. An implementation example of the `getProperties` method follows.

```

getProperties()

// getProperties method sample code
public Properties getProperties() {
    Properties p = new Properties();
    p.put("indexName", indexName);
    p.put("attributeName", attributeName);
    p.put("rangeIndex", ivRangeIndex ? "true" : "false");
    p.put("fieldAccessAttribute", ivFieldAccessAttribute ? "true" : "false");
    p.put("POJOKeyIndex", ivPOJOKeyIndex ? "true" : "false");
    return p;
}

```

setEntityMetadata method

The `setEntityMetadata` method is called by the WebSphere eXtreme Scale run time during initialization to set the `EntityMetadata` of the associated `BackingMap` on the `MapIndexPlugin` instance. The `EntityMetadata` is required for supporting indexing of tuple objects. A tuple is a data set that represents an entity object or its key. If the `BackingMap` is for an entity, then you must implement this method.

The following code sample implements the `setEntityMetadata` method.

```

setEntityMetadata(EntityMetadata entityMetadata)

// setEntityMetadata method sample code
public void setEntityMetadata(EntityMetadata entityMetadata) {
    ivEntityMetadata = entityMetadata;
    if (ivEntityMetadata != null) {
        // this is a tuple map
        TupleMetadata valueMetadata = ivEntityMetadata.getValueMetadata();
        int numAttributes = valueMetadata.getNumAttributes();
        for (int i = 0; i < numAttributes; i++) {
            String tupleAttributeName = valueMetadata.getAttribute(i).getName();
            if (attributeName.equals(tupleAttributeName)) {
                ivTupleValueIndex = i;
                break;
            }
        }
    }
}

```

```

    if (ivTupleValueIndex == -1) {
        // did not find the attribute in value tuple, try to find it on key tuple.
        // if found on key tuple, implies key indexing on one of tuple key attributes.
        TupleMetadata keyMetadata = ivEntityMetadata.getKeyMetadata();
        numAttributes = keyMetadata.getNumAttributes();
        for (int i = 0; i < numAttributes; i++) {
            String tupleAttributeName = keyMetadata.getAttribute(i).getName();
            if (attributeName.equals(tupleAttributeName)) {
                ivTupleValueIndex = i;
                ivKeyTupleAttributeIndex = true;
                break;
            }
        }
    }

    if (ivTupleValueIndex == -1) {
        // if entityMetadata is not null and we could not find the
        // attributeName in entityMetadata, this is an
        // error
        throw new ObjectGridRuntimeException("Invalid attributeName. Entity: " +
            ivEntityMetadata.getName());
    }
}
}

```

Attribute name methods

The `setAttributeName` method sets the name of the attribute to be indexed. The cached object class must provide the `get` method for the indexed attribute. For example, if the object has an `employeeName` or `EmployeeName` attribute, the index calls the `getEmployeeName` method on the object to extract the attribute value. The attribute name must be the same as the name in the `get` method, and the attribute must implement the `Comparable` interface. If the attribute is boolean type, you can also use the `isAttributeName` method pattern.

The `getAttributeName` method returns the name of the indexed attribute.

getAttribute method

The `getAttribute` method returns the indexed attribute value from the specified object. For example, if an `Employee` object has an attribute called `employeeName` that is indexed, you can use the `getAttribute` method to extract the `employeeName` attribute value from a specified `Employee` object. This method is required in a distributed WebSphere eXtreme Scale environment.

```

getAttribute(Object value)

// getAttribute method sample code
public Object getAttribute(Object value) throws ObjectGridRuntimeException {
    if (ivPOJOKeyIndex) {
        // In the POJO key indexing case, no need to get attribute from value object.
        // The key itself is the attribute value used to build the index.
        return null;
    }

    try {
        Object attribute = null;
        if (value != null) {
            // handle Tuple value if ivTupleValueIndex != -1
            if (ivTupleValueIndex == -1) {
                // regular value
                if (ivFieldAccessAttribute) {
                    attribute = this.getAttributeField(value).get(value);
                } else {
                    attribute = getAttributeMethod(value).invoke(value, emptyArray);
                }
            } else {
                // Tuple value
                attribute = extractValueFromTuple(value);
            }
        }
        return attribute;
    } catch (InvocationTargetException e) {
        throw new ObjectGridRuntimeException(
            "Caught unexpected Throwable during index update processing,
            index name = " + indexName + ": " + t,
            t);
    } catch (Throwable t) {

```

```

        throw new ObjectGridRuntimeException(
            "Caught unexpected Throwable during index update processing,
            index name = " + indexName + ": " + t,
            t);
    }
}

```

Plug-ins for indexing data

The built-in HashIndex, the `com.ibm.websphere.objectgrid.plugins.index.HashIndex` class, is a `MapIndexPlugin` plug-in that you can add into `BackingMap` to build static or dynamic indexes. This class supports both the `MapIndex` and `MapRangeIndex` interfaces. Defining and implementing indexes can significantly improve query performance.

Accessing data in WebSphere eXtreme Scale

After an application has a reference to an `ObjectGrid` instance or a client connection to a remote grid, you can access and interact with data in your WebSphere eXtreme Scale configuration. With the `ObjectGridManager` API, use one of the `createObjectGrid` methods to create a local instance, or the `getObjectGrid` method for a client instance with a distributed grid.

A thread in an application needs its own `Session`. When an application wants to use the `ObjectGrid` on a thread, it should just call one of the `getSession` methods to obtain a thread. This operation is cheap--there is no need to pool these operations in most cases. If the application is using a dependency injection framework such as Spring, you can inject a `Session` into an application bean when necessary.

After you obtain a `Session`, the application can access data stored in maps in the `ObjectGrid`. If the `ObjectGrid` uses entities, you can use the `EntityManager` API, which you can obtain with the `Session.getEntityManager` method. Because it is closer to Java specifications, the `EntityManager` interface is simpler than the map-based API. However, the `EntityManager` API carries a performance overhead because it tracks changes in objects. The map-based API is obtained by using the `Session.getMap` method.

WebSphere eXtreme Scale uses transactions. When an application interacts with a `Session`, it must be in the context of a transaction. A transaction is begun and committed or rolled back using the `Session.begin`, `Session.commit`, and `Session.rollback` methods on the `Session` object. Applications can also work in auto-commit mode, where the `Session` automatically begins and commits a transaction whenever the application interacts with `Maps`. However, the auto-commit mode is slower.

The logic of using transactions

Transactions may seem to be slow, but eXtreme Scale uses transactions for three reasons:

1. To allow rollback of changes if an exception occurs or business logic needs to undo state changes.
2. To hold locks on data and release locks within the lifetime of a transaction, allowing a set of changes to be made atomically, that is, all changes or no changes to data.
3. To produce an atomic unit of replication.

WebSphere eXtreme Scale lets a Session customize how much transaction is really needed. An application can turn off rollback support and locking but does so at a cost to the application. The application must handle the lack of these features itself.

For example, an application can turn off locking by configuring the BackingMap locking strategy to be NONE. This strategy is fast, but concurrent transactions can now modify the same data with no protection from each other. The application is responsible for all locking and data consistency when NONE is used.

An application can also change the way objects are copied when accessed by the transaction. The application can specify how objects are copied with the `ObjectMap.setCopyMode` method. With this method, you can turn off CopyMode. Turning off CopyMode is normally used for read-only transactions if different values can be returned for the same object within a transaction. Different values can be returned for the same object within a transaction.

For example, if the transaction called the `ObjectMap.get` method for the object at T1, it got the value at that point in time. If it calls the `get` method again within that transaction at a later time T2, another thread might have changed the value. Because the value has been changed by another thread, the application sees a different value. If the application modifies an object retrieved using a NONE CopyMode value, it is changing the committed copy of that object directly. Rolling back the transaction has no meaning in this mode. You are changing the only copy in the ObjectGrid. Although using the NONE CopyMode is fast, be aware of its consequences. An application that uses a NONE CopyMode must never roll back the transaction. If the application rolls back the transaction, the indexes are not updated with the changes *and* the changes are not replicated if replication is turned on. The default values are easy to use and less prone to errors. If you start trading performance in exchange for less reliable data, the application needs to be aware of what it is doing to avoid unintended problems.

CAUTION:

Be careful when you are changing either the locking or the CopyMode values. If you change the values, unpredictable application behavior will occur.

Interacting with stored data

After a session has been obtained, you can use the following code fragment to use the Map API for inserting data.

```
Session session = ...;
ObjectMap personMap = session.getMap("PERSON");
session.begin();
Person p = new Person();
p.name = "John Doe";
personMap.insert(p.name, p);
session.commit();
```

The same example using the EntityManager API follows. This code sample assumes that the Person object is mapped to an Entity.

```
Session session = ...;
EntityManager em = session.getEntityManager();
session.begin();
Person p = new Person();
p.name = "John Doe";
em.persist(p);
session.commit();
```

The pattern is designed to obtain references to the ObjectMaps for the Maps that the thread will work with, start a transaction, work with the data, then commit the transaction.

The ObjectMap interface has the typical Map operations such as put, get and remove. However, use the more specific operation names such as: get, getForUpdate, insert, update and remove. These method names convey the intent more precisely than the traditional Map APIs.

You can also use the indexing support, which is flexible.

The following is an example for updating an Object:

```
session.begin();
Person p = (Person)personMap.getForUpdate("John Doe");
p.name = "John Doe";
p.age = 30;
personMap.update(p.name, p);
session.commit();
```

The application normally uses the getForUpdate method rather than a simple get to lock the record. The update method must be called to actually provide the updated value to the Map. If update is not called then the Map is unchanged. The following is the same fragment using the EntityManager API:

```
session.begin();
Person p = (Person)em.findForUpdate(Person.class, "John Doe");
p.age = 30;
session.commit();
```

The EntityManager API is simpler than the Map approach. In this case, eXtreme Scale finds the Entity and returns a managed object to the application. The application modifies the object and commits the transaction, and eXtreme Scale tracks changes to managed objects automatically at commit time and performs the necessary updates.

Transactions and partitions

WebSphere eXtreme Scale transactions can only update a single partition. Transactions from a client can read from multiple partitions, but they can only update one partition. If an application attempts to update two partitions, then the transaction fails and is rolled back. A transaction that is using an embedded ObjectGrid (grid logic) has no routing capability and can only see data in the local partition. This business logic can always get a second session that is a true client session to access other partitions. However, this transaction would be an independent transaction.

Queries and partitions

If a transaction has already searched for an Entity, the transaction is associated with the partition for that Entity. Any queries that run on a transaction that is associated with an Entity are routed to the associated partition.

If a query is run on a transaction before it is associated with a partition, you must set the partition ID to use for the query. The partition ID is an integer value. The query is then routed to that partition.

Queries only search within a single partition. However, you can use the DataGrid APIs to run the same query in parallel on all partitions or a subset of partitions.

Use the DataGrid APIs to find an entry that might be in any partition.

7.0.0.0 FIX 2+ The REST data service allows any HTTP client to access a WebSphere eXtreme Scale grid, and is compatible with WCF Data Services in the Microsoft .NET Framework 3.5 SP1. For more information see the user guide for the eXtreme Scale REST data service

CopyMode best practices

WebSphere eXtreme Scale makes a copy of the value based on the six available CopyMode settings. Determine which setting works best for your deployment requirements.

You can use the BackingMap API `setCopyMode(CopyMode, valueInterfaceClass)` method to set the copy mode to one of the following final static fields that are defined in the `com.ibm.websphere.objectgrid.CopyMode` class.

When an application uses the ObjectMap interface to obtain a reference to a map entry, use that reference only within the WebSphere eXtreme Scale transaction that obtained the reference. Using the reference in a different transaction can lead to errors. For example, if you use the pessimistic locking strategy for the BackingMap, a `get` or `getForUpdate` method call acquires an S (shared) or U (update) lock, depending on the transaction. The `get` method returns the reference to the value and the lock that is obtained is released when the transaction completes. The transaction must call the `get` or `getForUpdate` method to lock the map entry in a different transaction. Each transaction must obtain its own reference to the value by calling the `get` or `getForUpdate` method instead of reusing the same value reference in multiple transactions.

CopyMode for entity maps

When using a map associated with an EntityManager API entity, the map always returns the entity Tuple objects directly without making a copy unless you are using COPY_TO_BYTES copy mode. It is important that the CopyMode is updated or the Tuple is copied appropriately when making changes.

COPY_ON_READ_AND_COMMIT

The COPY_ON_READ_AND_COMMIT mode is the default mode. The `valueInterfaceClass` argument is ignored when this mode is used. This mode ensures that an application does not contain a reference to the value object that is in the BackingMap. Instead, the application is always working with a copy of the value that is in the BackingMap. The COPY_ON_READ_AND_COMMIT mode ensures that the application can never inadvertently corrupt the data that is cached in the BackingMap. When an application transaction calls `ObjectMap.get` method for a given key, and it is the first access of the ObjectMap entry for that key, a copy of the value is returned. When the transaction is committed, any changes that are committed by the application are copied to the BackingMap to ensure that the application does not have a reference to the committed value in the BackingMap.

COPY_ON_READ

The `COPY_ON_READ` mode improves performance over the `COPY_ON_READ_AND_COMMIT` mode by eliminating the copy that occurs when a transaction is committed. The `valueInterfaceClass` argument is ignored when this mode is used. To preserve the integrity of the `BackingMap` data, the application ensures that every reference that it has for an entry is destroyed after the transaction is committed. With this mode, the `ObjectMap.get` method returns a copy of the value instead of a reference to the value to ensure that changes that are made by the application to the value does not affect the `BackingMap` value until the transaction is committed. However, when the transaction does commit, a copy of changes is not made. Instead, the reference to the copy that was returned by the `ObjectMap.get` method is stored in the `BackingMap`. The application destroys all map entry references after the transaction is committed. If application does not destroy the map entry references, the application might cause the data cached in `BackingMap` to become corrupted. If an application is using this mode and is having problems, switch to `COPY_ON_READ_AND_COMMIT` mode to see if the problem still exists. If the problem goes away, then the application is failing to destroy all of its references after the transaction has committed.

COPY_ON_WRITE

The `COPY_ON_WRITE` mode improves performance over the `COPY_ON_READ_AND_COMMIT` mode by eliminating the copy that occurs when the `ObjectMap.get` method is called for the first time by a transaction for a given key. The `ObjectMap.get` method returns a proxy to the value instead of a direct reference to the value object. The proxy ensures that a copy of the value is not made unless the application calls a set method on the value interface that is specified by the `valueInterfaceClass` argument. The proxy provides a copy on write implementation. When a transaction commits, the `BackingMap` examines the proxy to determine if any copy was made as a result of a set method being called. If a copy was made, then the reference to that copy is stored in the `BackingMap`. The big advantage of this mode is that a value is never copied on a read or at a commit when the transaction never calls a set method to change the value.

The `COPY_ON_READ_AND_COMMIT` and `COPY_ON_READ` modes both make a deep copy when a value is retrieved from the `ObjectMap`. If an application only updates some of the values that are retrieved in a transaction then this mode is not optimal. The `COPY_ON_WRITE` mode supports this behavior in an efficient manner but requires that the application uses a simple pattern. The value objects are required to support an interface. The application must use the methods on this interface when it is interacting with the value in an eXtreme Scale Session. If this is the case, then eXtreme Scale creates proxies for the values that are returned to the application. The proxy has a reference to the real value. If the application performs read operations only, the read operations always run against the real copy. If the application modifies an attribute on the object, the proxy makes a copy of the real object and then makes the modification on the copy. The proxy then uses the copy from that point on. Using the copy allows the copy operation to be avoided completely for objects that are only read by the application. All modify operations must start with the set prefix. Enterprise JavaBeans normally are coded to use this style of method naming for methods that modify the objects attributes. This convention must be followed. Any objects that are modified are copied at the time that they are modified by the application. This read and write scenario is the most efficient scenario supported by eXtreme Scale. To configure a map to use `COPY_ON_WRITE` mode, use the following example. In this example, the

application stores Person objects that are keyed using the name in the Map. The person object is represented in the following code snippet.

```
class Person {
    String name;
    int age;
    public Person() {
    }
    public void setName(String n) {
        name = n;
    }
    public String getName() {
        return name;
    }
    public void setAge(int a) {
        age = a;
    }
    public int getAge() {
        return age;
    }
}
```

The application uses the IPerson interface only when it interacts with values that are retrieved from a ObjectMap. Modify the object to use an interface as in the following example.

```
interface IPerson
{
    void setName(String n);
    String getName();
    void setAge(int a);
    int getAge();
}
// Modify Person to implement IPerson interface
class Person implements IPerson {
    ...
}
```

The application then needs to configure the BackingMap to use COPY_ON_WRITE mode, like in the following example:

```
ObjectGrid dg = ...;
BackingMap bm = dg.defineMap("PERSON");
// use COPY_ON_WRITE for this Map with
// IPerson as the valueProxyInfo Class
bm.setCopyMode(CopyMode.COPY_ON_WRITE,IPerson.class);
// The application should then use the following
// pattern when using the PERSON Map.
Session sess = ...;
ObjectMap person = sess.getMap("PERSON");
...
sess.begin();
// the application casts the returned value to IPerson and not Person
IPerson p = (IPerson)person.get("Billy");
p.setAge(p.getAge()+1);
...
// make a new Person and add to Map
Person p1 = new Person();
p1.setName("Bobby");
p1.setAge(12);
person.insert(p1.getName(), p1);
sess.commit();
// the following snippet WON'T WORK. Will result in ClassCastException
sess.begin();
```



```
// the mistake here is that Person is used rather than
// IPerson
Person a = (Person)person.get("Bobby");
sess.commit();
```

The first section shows the application retrieving a value that was named Billy in the map. The application casts the returned value to the IPerson object, not the Person object because the proxy that is returned implements two interfaces:

- The interface specified in the BackingMap.setCopyMode method call
- The com.ibm.websphere.objectgrid.ValueProxyInfo interface

You can cast the proxy to two types. The last part of the preceding code snippet demonstrates what is not allowed in COPY_ON_WRITE mode. The application retrieves the Bobby record and tries to cast the record to a Person object. This action fails with a class cast exception because the proxy that is returned is not a Person object. The returned proxy implements the IPerson object and ValueProxyInfo.

ValueProxyInfo interface and partial update support: This interface allows an application to retrieve either the committed read-only value referenced by the proxy or the set of attributes that have been modified during this transaction.

```
public interface ValueProxyInfo {
    List /**/ ibmGetDirtyAttributes();
    Object ibmGetRealValue();
}
```

The ibmGetRealValue method returns a read only copy of the object. The application must not modify this value. The ibmGetDirtyAttributes method returns a list of strings representing the attributes that have been modified by the application during this transaction. The main use case for ibmGetDirtyAttributes is in a Java database connectivity (JDBC) or CMP based loader. Only the attributes that are named in the list need be updated on either the SQL statement or object mapped to the table, which leads to more efficient SQL generated by the Loader. When a copy on write transaction is committed and if a loader is plugged in, the the loader can cast the values of the modified objects to the ValueProxyInfo interface to obtain this information.

Handling the equals method when using COPY_ON_WRITE or proxies: For example, the following code constructs a Person object and then inserts it to a an ObjectMap. Next, it retrieves the same object using the ObjectMap.get method. The value is cast to the interface. If the value is cast to the Person interface, a ClassCastException exception results because the returned value is a proxy that implements the IPerson interface and is not a Person object. The equality check fails when using the == operation because they are not the same object.

```
session.begin();
// new the Person object
Person p = new Person(...);
personMap.insert(p.getName, p);
// retrieve it again, remember to use the interface for the cast
IPerson p2 = personMap.get(p.getName());
if(p2 == p) {
    // they are the same
} else {
    // they are not
}
```

Another consideration is when you must override the equals method. As illustrated in the following snippet of code, the equals method must verify that the

argument is an object that implements IPerson interface and cast the argument to be a IPerson. Because the argument might be a proxy that implements the IPerson interface, you must use the getAge and getName methods when comparing instance variables for equality.

```
{
    if ( obj == null ) return false;
    if ( obj instanceof IPerson ) {
        IPerson x = (IPerson) obj;
        return ( age.equals( x.getAge() ) && name.equals( x.getName() ) )
    }
    return false;
}
```

ObjectQuery and HashIndex configuration requirements: When using COPY_ON_WRITE with ObjectQuery or a HashIndex plug-in, it's important to configure the ObjectQuery schema and HashIndex plug-in to access the objects using property methods, which is the default. If configured to use field access, the query engine and index will attempt to access the fields in the proxy object, which will always return null or 0 since the object instance will be a proxy.

NO_COPY

The NO_COPY mode allows an application to ensure that it never modifies a value object that is obtained using an ObjectMap.get method in exchange for performance improvements. The valueInterfaceClass argument is ignored when this mode is used. If this mode is used, no copy of the value is ever made. If the application modifies values, then the data in the BackingMap is corrupted. The NO_COPY mode is primarily useful for read-only maps where data is never modified by the application. If the application is using this mode and it is having problems, then switch to the COPY_ON_READ_AND_COMMIT mode to see if the problem still exists. If the problem goes away, then the application is modifying the value returned by ObjectMap.get method, either during transaction or after transaction has committed. All maps associated with EntityManager API entities automatically use this mode regardless of what is specified in the eXtreme Scale configuration.

All maps associated with EntityManager API entities automatically use this mode regardless of what is specified in the eXtreme Scale configuration.

COPY_TO_BYTES

You can store objects in a serialized format instead of POJO format. By using the COPY_TO_BYTES setting, you can reduce the memory footprint that a large graph of Objects can consume. See "Byte array maps" on page 41 for additional information.

Incorrect use of CopyMode

Errors occur when an application attempts to improve performance by using the COPY_ON_READ, COPY_ON_WRITE, or NO_COPY copy mode, as described above. The intermittent errors do not occur when you change the copy mode to the COPY_ON_READ_AND_COMMIT mode.

Problem

The problem might be due to corrupted data in the ObjectGrid map, which is a result of the application violating the programming contract of the copy mode that

is being used. Data corruption can cause unpredictable errors to occur intermittently or in an unexplained or unexpected fashion.

Solution

The application must comply with the programming contract that is stated for the copy mode being used. For the `COPY_ON_READ` and `COPY_ON_WRITE` copy modes, the application uses a reference to a value object outside of the transaction scope from which the value reference was obtained. To use these modes, the application must delete the reference to the value object after the transaction completes, and obtain a new reference to the value object in each transaction that accesses the value object. For the `NO_COPY` copy mode, the application must never change the value object. In this case, either write the application so that it does not change the value object, or set the application to use a different copy mode.

Byte array maps

You can store the key-value pairs in your maps in a byte array instead of POJO form, which reduces the memory footprint that a large graph of objects can consume.

Advantages

The amount of memory that is consumed increases with the number of objects in a graph of objects. By reducing a complicated graph of objects to a byte array, only one object is maintained in the heap instead of several objects. With this reduction of the number of objects in the heap, the Java run time has fewer objects to search for during garbage collection.

The default copy mechanism used by WebSphere eXtreme Scale is serialization, which is expensive. For instance, if using the default copy mode of `COPY_ON_READ_AND_COMMIT`, a copy is made both at read time and at get time. Instead of making a copy at read time, with byte arrays, the value is inflated from bytes, and instead of making a copy at commit time, the value is serialized to bytes. Using byte arrays results in equivalent data consistency to the default setting with a reduction of memory used.

When using byte arrays, note that having an optimized serialization mechanism is critical to seeing a reduction of memory consumption. For more information, see “Serialization performance” on page 214.

Configuring byte array maps

You can enable byte array maps with the ObjectGrid XML file by modifying the `CopyMode` attribute that is used by a map to the setting `COPY_TO_BYTES`, shown in the following example:

```
<backingMap name="byteMap" copyMode="COPY_TO_BYTES" />
```

See the topic on the ObjectGrid descriptor XML file in the *Administration Guide* for more information.

Considerations

You must consider whether or not to use byte array maps in a given scenario. Although you can reduce your memory use, processor use can increase when you use byte arrays.

The following list outlines several factors that should be considered before choosing to use the byte array map function.

Object type

Comparatively, memory reduction may not be possible when using byte array maps for some object types. Consequently, several types of objects exist for which you should not use byte array maps. If you are using any of the Java primitive wrappers as values, or a POJO that does not contain references to other objects (only storing primitive fields), the number of Java Objects is already as low as possible—there is only one. Since the amount of memory used by the object is already optimized, using a byte array map for these types of objects is not recommended. Byte array maps are more suitable to object types that contain other objects or collections of objects where the total number of POJO objects is greater than one.

For example, if you have a Customer object that had a business Address and a home Address, as well as a collection of Orders, the number of objects in the heap and the number of bytes used by those objects can be reduced by using byte array maps.

Local access

When using other copy modes, applications can be optimized when copies are made if objects are Cloneable with the default ObjectTransformer or when a custom ObjectTransformer is provided with an optimized copyValue method. Compared to the other copy modes, copying on reads, writes, or commit operations will have additional cost when accessing objects locally. For example, if you have a near cache in a distributed topology or are directly accessing a local or server ObjectGrid instance, the access and commit time will increase when using byte array maps due to the cost of serialization. You will see a similar cost in a distributed topology if you use data grid agents or you access the server primary when using the ObjectGridEventGroup.ShardEvents plug-in.

Plug-in interactions

With byte array maps, objects are not inflated when communicating from a client to a server unless the server needs the POJO form. Plug-ins that interact with the map value will experience a reduction in performance due to the requirement to inflate the value.

Any plug-in that uses LogElement.getCacheEntry or LogElement.getCurrentValue will see this additional cost. If you want to get the key, you can use LogElement.getKey, which avoids the additional overhead associated with the LogElement.getCacheEntry().getKey method. The following sections discuss plug-ins in light of the usage of byte arrays.

Indexes and queries

When objects are stored in POJO format, the cost of doing indexing and querying is minimal because the object does not need to be inflated. When using a byte array map you will have the additional cost of inflating the object. In general if your application uses indexes or queries, it is not recommended to use byte array maps unless you only run queries on key attributes.

Optimistic locking

When using the optimistic locking strategy, you will have the additional cost during updates and invalidate operations. This comes from having to inflate the value on the server to get the version value to do optimistic collision checking. If you are just using optimistic locking to guarantee fetch operations and do not need optimistic collision checking, you can use the `com.ibm.websphere.objectgrid.plugins.builtins.NoVersioningOptimisticCallback` to disable version checking.

Loader

With a Loader, you will also have the cost in the eXtreme Scale run time from inflating and reserializing the value when it is used by the Loader. You can still use byte array maps with Loaders, but consider the cost of making changes to the value in such a scenario. For example, you can use the byte array feature in the context of a read mostly cache. In this case, the benefit of having less objects in the heap and less memory used will outweigh the cost incurred from using byte arrays on insert and update operations.

ObjectGridEventListener

When using the `transactionEnd` method in the `ObjectGridEventListener` plug-in, you will have an additional cost on the server side for remote requests when accessing a `LogElement`'s `CacheEntry` or current value. If the implementation of the method does not access these fields, then you will not have the additional cost.

Using Sessions to access data in the grid

Applications can begin and end transactions through the `Session` interface. The `Session` interface also provides access to the application-based `ObjectMap` and `JavaMap` interfaces.

Each `ObjectMap` or `JavaMap` instance is directly tied to a specific `Session` object. Each thread that wants access to an eXtreme Scale must first obtain a `Session` from the `ObjectGrid` object. A `Session` instance cannot be shared concurrently between threads. WebSphere eXtreme Scale does not use any thread local storage, but platform restrictions might limit the opportunity to pass a `Session` from one thread to another.

Methods

Get method

An application obtains a `Session` instance from an `ObjectGrid` object using the `ObjectGrid.getSession` method. The following example demonstrates how to obtain a `Session` instance:

```
ObjectGrid objectGrid = ...; Session sess = objectGrid.getSession();
```

After a Session is obtained, the thread keeps a reference to the session for its own use. Calling the getSession method multiple times returns a new Session object each time.

Transactions and Session methods

A Session can be used to begin, commit, or rollback transactions. Operations against BackingMaps using ObjectMaps and JavaMaps are most efficiently performed within a Session transaction. After a transaction has started, any changes to one or more BackingMaps in that transaction scope are stored in a special transaction cache until the transaction is committed. When a transaction is committed, the pending changes are applied to the BackingMaps and Loaders and become visible to any other clients of that ObjectGrid.

WebSphere eXtreme Scale also supports the ability to automatically commit transactions, also known as auto-commit. If any ObjectMap operations are performed outside of the context of an active transaction, an implicit transaction is started before the operation and the transaction is automatically committed before returning control to the application.

```
Session session = objectGrid.getSession();
ObjectMap objectMap = session.getMap("someMap");
session.begin();
objectMap.insert("key1", "value1");
objectMap.insert("key2", "value2");
session.commit();
objectMap.insert("key3", "value3"); // auto-commit
```

Session.flush method

The Session.flush method only makes sense when a Loader is associated with a BackingMap. The flush method invokes the Loader with the current set of changes in the transaction cache. The Loader applies the changes to the backend. These changes are not committed when the flush is invoked. If a Session transaction is committed after a flush invocation, only updates that happen after the flush invocation are applied to the Loader. If a Session transaction is rolled back after a flush invocation, the flushed changes are discarded with all other pending changes in the transaction. Use the Flush method sparingly because it limits the opportunity for batch operations against a Loader. Following is an example of the usage of the Session.flush method:

```
Session session = objectGrid.getSession();
session.begin();
// make some changes
...
session.flush(); // push these changes to the Loader, but don't commit yet
// make some more changes
...
session.commit();
```

NoWriteThrough method

Some maps are backed by a Loader, which provides persistent storage for the data in the map. Sometimes it is useful to commit data just to the eXtreme Scale map and not push data out to the Loader. The Session interface provides the beginNoWriteThrough method for this purpose. The beginNoWriteThrough method starts a transaction like the begin method. With the beginNoWriteThrough method, when the transaction is committed, the data is only committed to the in-memory map and is not committed to the persistent storage that is provided by the Loader. This method is very useful when performing data preload on the map.

When using a distributed ObjectGrid instance, the `beginNoWriteThrough` method is useful for making changes to the near cache only, without modifying the far cache on the server. If the data is known to be stale in the near cache, using the `beginNoWriteThrough` method can allow entries to be invalidated on the near cache without invalidating them on the server as well.

The Session interface also provides the `isWriteThroughEnabled` method to determine what type of transaction is currently active.

```
Session session = objectGrid.getSession();
session.beginNoWriteThrough();
// make some changes ...
session.commit(); // these changes will not get pushed to the Loader
```

Obtain the TxID object method

The TxID object is an opaque object that identifies the active transaction. Use the TxID object for the following purposes:

- For comparison when you are looking for a particular transaction.
- To store shared data between the TransactionCallback and Loader objects.

See TransactionCallback plug-in and Loaders for additional information about the Object slot feature.

Performance monitoring method

If you are using eXtreme Scale within WebSphere Application Server, it might be necessary to reset the transaction type for performance monitoring. You can set the transaction type with the `setTransactionType` method. See Monitoring ObjectGrid performance with WebSphere Application Server performance monitoring infrastructure (PMI) for more information about the `setTransactionType` method.

Process a complete LogSequence method

WebSphere eXtreme Scale can propagate sets of map changes to ObjectGrid listeners as a means of distributing maps from one Java virtual machine to another. To make it easier for the listener to process the received LogSequences, the Session interface provides the `processLogSequence` method. This method examines each LogElement within the LogSequence and performs the appropriate operation, for example, insert, update, invalidate, and so on, against the BackingMap that is identified by the LogSequence MapName. An ObjectGrid Session must be available before the `processLogSequence` method is invoked. The application is also responsible for issuing the appropriate commit or rollback calls to complete the Session. Autocommit processing is not available for this method invocation. Normal processing by the receiving ObjectGridEventListener at the remote JVM would be to start a Session using the `beginNoWriteThrough` method, which prevents endless propagation of changes, followed by a call to this `processLogSequence` method, and then committing or rolling back the transaction.

```
// Use the Session object that was passed in during
//ObjectGridEventListener.initialization...
session.beginNoWriteThrough();
// process the received LogSequence
try {
    session.processLogSequence(receivedLogSequence);
} catch (Exception e) {
```

```
    session.rollback(); throw e;
}
// commit the changes
session.commit();
```

markRollbackOnly method

This method is used to mark the current transaction as "rollback only". Marking a transaction "rollback only" ensures that even if the commit method is called by application, the transaction is rolled back. This method is typically used by ObjectGrid itself or by the application when it knows that data corruption could occur if the transaction was allowed to be committed. After this method is called, the Throwable object that is passed to this method is chained to the `com.ibm.websphere.objectgrid.TransactionException` exception that results by the commit method if it is called on a Session that was previously marked a "rollback only". Any subsequent calls to this method for a transaction that is already marked as "rollback only" is ignored. That is, only the first call that passes a non-null Throwable reference is used. Once the marked transaction is completed, the "rollback only" mark is removed so that the next transaction that is started by the Session can be committed.

isMarkedRollbackOnly method

Returns if Session is currently marked as "rollback only". Boolean true is returned by this method if and only if `markRollbackOnly` method was previously called on this Session and the transaction started by the Session is still active.

setTransactionTimeout method

Set transaction timeout for next transaction started by this Session to a specified number of seconds. This method does not affect the transaction timeout of any transactions previously started by this Session. It only affects transactions that are started after this method is called. If this method is never called, then the timeout value that was passed to the `setTxTimeout` method of the `com.ibm.websphere.objectgrid.ObjectGrid` method is used.

getTransactionTimeout method

This method returns the transaction timeout value in seconds. The last value that was passed as the timeout value to the `setTransactionTimeout` method is returned by this method. If the `setTransactionTimeout` method is never called, then the timeout value that was passed to the `setTxTimeout` method of the `com.ibm.websphere.objectgrid.ObjectGrid` method is used.

transactionTimedOut

This method returns boolean true if the current transaction that was started by this Session has timed out.

isFlushing method

This method returns boolean true if and only if all transaction changes are being flushed out to the Loader plug-in as a result of the `flush` method of Session interface being invoked. A Loader plug-in may find this method useful when it needs to know why its `batchUpdate` method was invoked.

isCommitting method

This method returns boolean true if and only if all transaction changes are being committed as a result of the commit method of Session interface being invoked. A Loader plug-in might find this method useful when it needs to know why its batchUpdate method was invoked.

setRequestRetryTimeout method

This method sets the request retry timeout value for the session in milliseconds. If the client set a request retry timeout, the session setting overrides the client value.

getRequestRetryTimeout method

This method gets the current request retry timeout setting on the session. A value of -1 indicates that the timeout is not set. A value of 0 indicates it is in fail-fast mode. A value greater than 0 indicates the timeout setting in milliseconds.

SessionHandle for routing

When you are using a per-container partition placement policy, you can use a SessionHandle object. A SessionHandle object contains partition information for the current Session and can be reused for a new Session.

A SessionHandle object includes information for the partition to which the current Session is bound. SessionHandle is extremely useful for the per-container partition placement policy and can be serialized with standard Java serialization.

If you have a SessionHandle object, you can apply that handle to a Session with the setSessionHandle(SessionHandle target) method, passing the handle in as the target. You can retrieve the SessionHandle object with the Session.getSessionHandle method.

Because it is only applicable in a per-container placement scenario, getting the SessionHandle object throws an IllegalStateException if a given data grid has multiple per-container map sets or has no per-container map sets. If you do not invoke the setSessionHandle method before calling the getSessionHandle method, the appropriate SessionHandle object is selected based on your client properties configuration.

You can also use the SessionHandleTransformer helper class to convert the handle into different formats. The methods of this class can change a handle's representation from byte array to instance, string to instance, and vice versa for both cases, and can also write the handle's contents into the output stream.

For an example on how you can use a SessionHandle object, see the zone-preferred routing topic in the .

SessionHandle integration

A SessionHandle object includes the partition information for the Session to which it is bound and facilitates request routing. SessionHandle objects apply to the per-container partition placement scenario only.

SessionHandle object for request routing

You can bind a SessionHandle object to a Session in the following ways:

Tip: In each of the following method calls, after a SessionHandle object is bound to a Session, the SessionHandle object can be obtained from the Session.getSessionHandle method.

- **Call the Session.getSessionHandle method:** When this method is called, if no SessionHandle object is bound to the Session, a SessionHandle object is selected randomly and bound to the Session.
- **Call transactional create, read, update, delete operations:** When these methods are called or at commit time, if no SessionHandle object is bound to the Session, a SessionHandle object is selected randomly and bound to the Session.
- **Call ObjectMap.getNextKey method:** When this method is called, if no SessionHandle object is bound to the Session, the operation request is randomly routed to individual partitions until a key is obtained. If a key is returned from a partition, a SessionHandle object that corresponds to that partition is bound to the Session. If no key is found, no SessionHandle is bound to the Session.
- **Call the QueryQueue.getNextEntity or QueryQueue.getNextEntities methods:** At the time this method is called, if no SessionHandle object is bound to the Session, the operation request is randomly routed to individual partitions until an object is obtained. If an object is returned from a partition, a SessionHandle object that corresponds to that partition is bound to the Session. If no object is found, no SessionHandle is bound to the Session.
- **Set a SessionHandle with the Session.setSessionHandle(SessionHandle sh) method:** If a SessionHandle is obtained from the Session.getSessionHandle method, the SessionHandle can be bound to a Session. Setting a SessionHandle influences request routing within the scope of the Session to which it is bound.

The Session.getSessionHandle method always returns a SessionHandle object. The method also automatically binds a SessionHandle on the Session if no SessionHandle object is bound to the Session. If you want to verify whether a Session has a SessionHandle object only, call the Session.isSessionHandleSet method. If this method returns a value of false, no SessionHandle object is currently bound to the Session.

Major operation types in the per-container placement scenario

A summary of the routing behavior of major operation types in the per-container partition placement scenario with respect to SessionHandle objects follows.

- **Session object with bound SessionHandle object**
 - Index - MapIndex and MapRangeIndex API: SessionHandle
 - Query and ObjectQuery: SessionHandle
 - Agent - MapGridAgent and ReduceGridAgent API: SessionHandle
 - ObjectMap.Clear: SessionHandle
 - ObjectMap.getNextKey: SessionHandle
 - QueryQueue.getNextEntity, QueryQueue.getNextEntities: SessionHandle
 - Transactional create, retrieve, update, and delete operations (ObjectMap API and EntityManager API): SessionHandle
- **Session object without bound SessionHandle object**
 - Index - MapIndex and MapRangeIndex API: All current active partitions
 - Query and ObjectQuery: Specified partition with setPartition method of Query and ObjectQuery
 - Agent - MapGridAgent and ReduceGridAgent

- Not supported: `ReduceGridAgent.reduce(Session s, ObjectMap map, Collection keys)` and `MapGridAgent.process(Session s, ObjectMap map, Object key)` method.
- All current active partitions: `ReduceGridAgent.reduce(Session s, ObjectMap map)` and `MapGridAgent.processAllEntries(Session s, ObjectMap map)` method.
- `ObjectMap.clear`: All current active partitions.
- `ObjectMap.getNextKey`: Binds a `SessionHandle` to the `Session` if a key is returned from one of the randomly selected partitions. If no key is returned, the `Session` is not bound to a `SessionHandle`.
- `QueryQueue`: Specifies a partition with the `QueryQueue.setPartition` method. If no partition is set, the method randomly selects a partition to return. If an object is returned, the current `Session` is bound with the `SessionHandle` that is bound to the partition that returns the object. If no object is returned, the `Session` is not bound to a `SessionHandle`.
- Transactional create, retrieve, update, and delete operations (`ObjectMap` API and `EntityManager` API): Randomly select a partition.

In most cases, use `SessionHandle` to control routing to a particular partition. You can retrieve and cache the `SessionHandle` from the `Session` that inserts data. After caching the `SessionHandle`, you can set it on another `Session` so that you can route requests to the partition specified by the cached `SessionHandle`. To perform operations such as `ObjectMap.clear` without `SessionHandle`, you can temporarily set the `SessionHandle` to null by calling `Session.setSessionHandle(null)`. Without a specified `SessionHandle`, operations run on all current active partitions.

- **QueryQueue routing behavior**

In the per-container partition placement scenario, you can use `SessionHandle` to control routing of `getNextEntity` and `getNextEntities` methods of the `QueryQueue` API. If the `Session` is bound to a `SessionHandle`, requests route to the partition to which the `SessionHandle` is bound. If the `Session` is not bound to a `SessionHandle`, requests are routed to the partition set with the `QueryQueue.setPartition` method if a partition has been set in this way. If the `Session` has no bound `SessionHandle` or partition, a randomly selected partition are returned. If no such partition is found, the process stops and no `SessionHandle` is bound to the current `Session`.

The following snippet of code shows how to use `SessionHandle` objects.

```
Session ogSession = objectGrid.getSession();

// binding the SessionHandle
SessionHandle sessionHandle = ogSession.getSessionHandle();

ogSession.begin();
ObjectMap map = ogSession.getMap("planet");
map.insert("planet1", "mercury");

// transaction is routed to partition specified by SessionHandle
ogSession.commit();

// cache the SessionHandle that inserts data
SessionHandle cachedSessionHandle = ogSession.getSessionHandle();

// verify if SessionHandle is set on the Session
boolean isSessionHandleSet = ogSession.isSessionHandleSet();

// temporarily unbind the SessionHandle from the Session
if(isSessionHandleSet) {
```

```

        ogSession.setSessionHandle(null);
    }

    // if the Session has no SessionHandle bound,
    // the clear operation will run on all current active partitions
    // and thus remove all data from the map in the entire grid
    map.clear();

    // after clear is done, reset the SessionHandle back,
    // if the Session needs to use previous SessionHandle.
    // Optionally, calling getSessionHandle can get a new SessionHandle
    ogSession.setSessionHandle(cachedSessionHandle);

```

Application design considerations

In the per-container placement strategy scenario, use the SessionHandle object for most operations. The SessionHandle object controls routing to partitions. To retrieve data, the SessionHandle object that you bind to the Session must be same SessionHandle object from any insert data transaction.

When you want to perform an operation without a SessionHandle set on the Session, you can unbind a SessionHandle from a Session by making a Session.setSessionHandle(null) method call.

When a Session is bound to a SessionHandle, all operation requests route to the partition that is specified by the SessionHandle object. Without the SessionHandle object set, operations route to either all partitions or a randomly selected partition.

Caching objects with no relationships involved (ObjectMap API)

ObjectMaps are like Java Maps that allow data to be stored as key-value pairs. ObjectMaps provide a simple and intuitive approach for the application to store data. An ObjectMap is ideal for caching objects that have no relationships involved. If object relationships are involved, then you should use the EntityManager API.

For more information about the EntityManager API, see “Caching objects and their relationships (EntityManager API)” on page 60.

Applications typically obtain a WebSphere eXtreme Scale reference and then obtain a Session object from the reference for each thread. Sessions cannot be shared between threads. The getMap method of Session returns a reference to an ObjectMap to use for this thread.

Introduction to ObjectMap

The ObjectMap interface is used for transactional interaction between applications and BackingMaps.

Purpose

An ObjectMap instance is obtained from a Session object that corresponds to the current thread. The ObjectMap interface is the main vehicle that applications use to make changes to entries in a BackingMap.

Obtain an ObjectMap instance

An application gets an ObjectMap instance from a Session object using the Session.getMap(String) method. The following code snippet demonstrates how to obtain an ObjectMap instance:

```
ObjectGrid objectGrid = ...;
BackingMap backingMap = objectGrid.defineMap("mapA");
Session sess = objectGrid.getSession();
ObjectMap objectMap = sess.getMap("mapA");
```

Each ObjectMap instance corresponds to a particular Session object. Calling the getMap method multiple times on a particular Session object with the same BackingMap name always returns the same ObjectMap instance.

Automatically commit transactions

Operations against BackingMaps that use ObjectMaps and JavaMaps are performed most efficiently within a Session transaction. WebSphere eXtreme Scale provides autocommit support when methods on the ObjectMap and JavaMap interfaces are called outside of a Session transaction. The methods start an implicit transaction, perform the requested operation, and commit the implicit transaction.

Method semantics

An explanation of the semantics behind each method on the ObjectMap and JavaMap interfaces follows. The setDefaultKeyword method, the invalidateUsingKeyword method, and the methods that have a Serializable argument are discussed in the Keywords topic. The setTimeToLive method is discussed in the Evictors topic. See the API documentation for more information on these methods.

containsKey method

The containsKey method determines if a key has a value in the BackingMap or Loader. If null values are supported by an application, this method can be used to determine if a null reference that is returned from a get operation refers to a null value or indicates that the BackingMap and Loader do not contain the key.

flush method

The flush method semantics are similar to the flush method on the Session interface. The notable difference is that the Session flush applies the current pending changes for all of the maps that are modified in the current session. With this method, only the changes in this ObjectMap instance are flushed to the loader.

get method

The get method fetches the entry from the BackingMap instance. If the entry is not found in the BackingMap instance but a Loader is associated with the BackingMap instance, the BackingMap instance attempts to fetch the entry from the Loader. The getAll method is provided to allow batch fetch processing.

getForUpdate method

The getForUpdate method is the same as the get method, but using the getForUpdate method tells the BackingMap and Loader that the intention is to update the entry. A Loader can use this hint to issue a SELECT for UPDATE query to a database backend. If a pessimistic locking strategy is

defined for the BackingMap, the lock manager locks the entry. The `getAllForUpdate` method is provided to allow batch fetch processing.

insert method

The `insert` method inserts an entry into the BackingMap and the Loader. Using this method tells the BackingMap and Loader that you want to insert an entry that did not previously exist. When you invoke this method on an existing entry, an exception occurs when the method is invoked or when the current transaction is committed.

invalidate method

The semantics of the `invalidate` method depend on the value of the `isGlobal` parameter that is passed to the method. The `invalidateAll` method is provided to allow batch invalidate processing.

Local invalidation is specified when the value `false` is passed as the `isGlobal` parameter of the `invalidate` method. Local invalidation discards any changes to the entry in the transaction cache. If the application issues a `get` method, the entry is fetched from the last committed value in the BackingMap. If no entry is present in the BackingMap, the entry is fetched from the last flushed or committed value in the Loader. When a transaction is committed, any entries that are marked as locally invalidated have no impact on the BackingMap. Any changes that were flushed to the Loader are still committed even if the entry was invalidated.

Global invalidation is specified when `true` is passed as the `isGlobal` parameter of the `invalidate` method. Global invalidation discards any pending changes to the entry in the transaction cache and bypasses the BackingMap value on subsequent operations that are performed on the entry. When a transaction is committed, any entries that are marked as globally invalidated are evicted from the BackingMap. Consider the following use case for invalidation as an example: The BackingMap is backed by a database table that has an auto increment column. Increment columns are useful for assigning unique numbers to records. The application inserts an entry. After the insert, the application needs to know the sequence number for the inserted row. It knows that its copy of the object is old, so it uses global invalidation to get the value from the Loader. The following code demonstrates this use case:

```
Session sess = objectGrid.getSession();
ObjectMap map = sess.getMap("mymap");
sess.begin();
map.insert("Billy", new Person("Joe", "Bloggs", "Manhattan"));
sess.flush();
map.invalidate("Billy", true);
Person p = map.get("Billy");
System.out.println("Version column is: " + p.getVersion());
map.commit();
```

This code sample adds an entry for Billy. The version attribute of `Person` is set using an auto-increment column in the database. The application first performs an insert command. It then issues a flush, which causes the insert to be sent to the Loader and database. The database sets the version column to the next number in the sequence, which makes the `Person` object in the transaction outdated. To update the object, the application is globally invalidated. The next `get` method that is issued gets the entry from the Loader, ignoring the transaction value. The entry is fetched from the database with the updated version value.

put method

The semantics of the `put` method are dependent on whether a previous `get`

method was invoked in the transaction for the key. If the application issues a get operation that returns an entry that exists in the BackingMap or Loader, the put method invocation is interpreted as an update and returns the previous value in the transaction. If a put method invocation ran without a previous get method invocation, or a previous get method invocation did not find an entry, the operation is interpreted as an insert. The semantics of the insert and update methods apply when the put operation is committed. The putAll method is provided to enable batch insert and update processing.

remove method

The remove method removes the entry from the BackingMap and the Loader, if a Loader is plugged in. The value of the object that was removed is returned by this method. If the object does not exist, this method returns a null value. The removeAll method is provided to enable batch deletion processing without the return values.

setCopyMode method

The setCopyMode method specifies a CopyMode value for this ObjectMap. With this method, an application can override the CopyMode value that is specified on the BackingMap. The specified CopyMode value is in effect until clearCopyMode method is invoked. Both methods are invoked outside of transactional bounds. A CopyMode value cannot be changed in the middle of a transaction.

touch method

The touch method updates the last access time for an entry. This method does not retrieve the value from the BackingMap. Use this method in its own transaction. If the provided key does not exist in the BackingMap because of invalidation or removal, an exception occurs during commit processing.

update method

The update method explicitly updates an entry in the BackingMap and the Loader. Using this method indicates to the BackingMap and Loader that you want to update an existing entry. An exception occurs if you invoke this method on an entry that does not exist when the method is invoked or during commit processing.

getIndex method

The getIndex method attempts to obtain a named index that is built on the BackingMap. The index cannot be shared between threads and works on the same rules as a Session. The returned index object should be cast to the right application index interface such as the MapIndex interface, the MapRangeIndex interface, or a custom index interface.

clear method

The clear method removes all cache entries from a map from all partitions. This operation is an auto-commit function, so no active transaction should be present when calling clear.

Note: The clear method only clears out the map on which it is called, leaving any related entity maps unaffected. This method does not invoke the Loader plug-in.

Dynamic maps

With dynamic maps, you can create maps after the data grid has already been initialized.

In previous versions, WebSphere eXtreme Scale has required you to define maps before initializing the ObjectGrid. As a result, you had to create all of the maps to be used before running transactions against any of the maps.

Advantages of dynamic maps

The introduction of dynamic maps reduces the restriction of having to define all maps before initialization. Through the use of template maps, you can create maps after the ObjectGrid has been initialized.

Template maps are defined in the ObjectGrid XML file. Template comparisons are run when a Session requests a map that has not been previously defined. If the new map name matches the regular expression of a template map, the map is created dynamically and assigned the name of the requested map. This newly created map inherits all of the settings of the template map as defined by the ObjectGrid XML file.

Creating dynamic maps

Dynamic map creation is tied to the `Session.getMap(String)` method. Calls to this method return an `ObjectMap` based on the `BackingMap` that was configured by the ObjectGrid XML file.

Passing in a `String` that matches the regular expression of a template map results in the creation of an `ObjectMap` and an associated `BackingMap`.

See the API documentation for more information about the `Session.getMap(String cacheName)` method.

Defining a template map in XML is as simple as setting a template Boolean attribute on the `BackingMap` element. When `template` is set to `true`, the name of the `BackingMap` is interpreted as a regular expression.

WebSphere eXtreme Scale uses Java regular expression pattern matching. For more information about the regular expression engine in Java, see the API documentation for the `java.util.regex` package and classes.

Note: When you are defining template maps, verify that the map names are unique enough so that the application can match to only one of the template maps with the `Session.getMap(String mapName)` method. If the `getMap()` method matches more than one template map pattern, an `IllegalArgumentException` exception results.

A sample ObjectGrid XML file with a template map defined follows.

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">
  <objectGrids>
    <objectGrid name="accounting">
      <BackingMap name="payroll" readOnly="false" />
      <BackingMap name="templateMap.*" template="true"
        pluginCollectionRef="templatePlugins" lockStrategy="PESSIMISTIC" />
    </objectGrid>
  </objectGrids>

  <BackingMapPluginCollections>
    <BackingMapPluginCollection id="templatePlugins">
```



```

    <bean id="Evictor"
        className="com.ibm.websphere.objectgrid.plugins.builtins.LFUEvictor" />
    </backingMapPluginCollection>
</backingMapPluginCollections>
</objectGridConfig>

```

The previous XML file defines one template map and one non-template map. The name of the template map is a regular expression: `templateMap.*`. When the `Session.getMap(String)` method is called with a map name matching this regular expression, the application creates a map.

Example

Configuration of a template map is required in order to create a dynamic map. Add the template Boolean to a backingMap in the ObjectGrid XML file.

```
<backingMap name="templateMap.*" template="true" />
```

The name of the template map is treated as a regular expression.

Calling the `Session.getMap(String cacheName)` method with a `cacheName` that is a match for the regular expression results in the creation of the dynamic map. An `ObjectMap` object is returned from this method call, and an associated `BackingMap` object is created.

```

Session session = og.getSession();
ObjectMap map = session.getMap("templateMap1");

```

The newly created map is configured with all the attributes and plug-ins that were defined on the template map definition. Consider again the previous ObjectGrid XML file.

A dynamic map created based on the template map in this XML file would have an evictor configured and its lock strategy would be pessimistic.

Note: A template is not an actual `BackingMap`. That is, the “accounting” ObjectGrid does not contain an actual `“templateMap.*”` map. The template is only used as a basis for dynamic map creation. However, you must include the dynamic map in the `mapRef` element of the deployment policy XML file named exactly as in the ObjectGrid XML. This element identifies which `mapSet` in which the dynamic maps are defined.

Consider the change in behavior of the `Session.getMap(String cacheName)` method when using template maps. Before WebSphere eXtreme Scale Version 7.0, all calls to the `Session.getMap(String cacheName)` method resulted in an `UndefinedMapException` exception if the map requested did not exist. With dynamic maps, every name that matches the regular expression for a template map results in map creation. Be sure to note the number of maps that your application creates, particularly if your regular expression is generic.

Also, `ObjectGridPermission.DYNAMIC_MAP` is required for dynamic map creation when eXtreme Scale security is enabled. This permission is checked when the `Session.getMap(String)` method is called. For more information, see the information about application client authorization in the *Product Overview*.

Additional examples

objectGrid.xml

```

<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">
<objectGrids>
  <objectGrid name="session">
    <backingMap name="objectgrid.session.metadata.dynamicmap.*" template="true"
      lockStrategy="PESSIMISTIC" ttlEvictorType="LAST_ACCESS_TIME">
    <backingMap name="objectgrid.session.attribute.dynamicmap.*"
      template="true" lockStrategy="OPTIMISTIC"/>
    <backingMap name="datagrid.session.global.ids.dynamicmap.*"
      lockStrategy="PESSIMISTIC"/>
  </objectGrid>
</objectGrids>
</objectGridConfig>

```

objectGridDeployment.xml

```

<?xml version="1.0" encoding="UTF-8"?>
<deploymentPolicy xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/ws/objectgrid/deploymentPolicy
../deploymentPolicy.xsd"
  xmlns="http://ibm.com/ws/objectgrid/deploymentPolicy">
<objectgridDeployment objectgridName="session">
  <mapSet name="mapSet2" numberOfPartitions="5" minSyncReplicas="0"
    maxSyncReplicas="0" maxAsyncReplicas="1" developmentMode="false"
    placementStrategy="PER_CONTAINER">
    <map ref="logical.name"/>
    <map ref="objectgrid.session.metadata.dynamicmap.*"/>
    <map ref="objectgrid.session.attribute.dynamicmap.*"/>
    <map ref="datagrid.session.global.ids"/>
  </mapSet>
</objectgridDeployment>
</deploymentPolicy>

```

Limitations and considerations

Limitations:

- The QuerySchema element does not support the template for mapName.
- You cannot use entities with dynamic maps.
- An entity BackingMap is implicitly defined, mapped to the entity through the class name.

Considerations:

- Many plug-ins have no way of determining the map with which each plug-in is associated.
- Other plug-ins differentiate themselves by using a map name or BackingMap name as an argument.
- When you are defining template maps, verify that the map names are unique enough so that the application can match to only one of the template maps using the Session.getMap(String mapName) method. If the getMap() method matches more than one template map pattern, an IllegalArgumentException exception results.

ObjectMap and JavaMap

A JavaMap instance is obtained from an ObjectMap object. The JavaMap interface has the same method signatures as ObjectMap, but with different exception handling. JavaMap extends the java.util.Map interface, so all exceptions are instances of the java.lang.RuntimeException class. Because JavaMap extends the

java.util.Map interface, it is easy to quickly use WebSphere eXtreme Scale with an existing application that uses a java.util.Map interface for object caching.

Obtain a JavaMap instance

An application gets a JavaMap instance from an ObjectMap object using the ObjectMap.getJavaMap method. The following code snippet demonstrates how to obtain a JavaMap instance.

```
ObjectGrid objectGrid = ...;
BackingMap backingMap = objectGrid.defineMap("mapA");
Session sess = objectGrid.getSession();
ObjectMap objectMap = sess.getMap("mapA");
java.util.Map map = objectMap.getJavaMap();
JavaMap javaMap = (JavaMap) javaMap;
```

A JavaMap is backed by the ObjectMap from which it was obtained. Calling the getJavaMap method multiple times using a particular ObjectMap always returns the same JavaMap instance.

Methods

The JavaMap interface only supports a subset of the methods on the java.util.Map interface. The java.util.Map interface supports the following methods:

containsKey(java.lang.Object) method

get(java.lang.Object) method

put(java.lang.Object, java.lang.Object) method

putAll(java.util.Map) method

remove(java.lang.Object) method

clear()

All other methods inherited from the java.util.Map interface result in a java.lang.UnsupportedOperationException exception.

Maps as FIFO queues

With WebSphere eXtreme Scale, you can provide a first-in first-out (FIFO) queue-like capability for all maps. WebSphere eXtreme Scale tracks the insertion order for all maps. A client can ask a map for the next unlocked entry in a map in the order of insertion and lock the entry. This process allows multiple clients to consume entries from the map efficiently.

FIFO example

The following code snippet shows a client entering a loop to process entries from the map until the map is exhausted. The loop starts a transaction, then calls the ObjectMap.getNextKey(5000) method. This method returns the key of the next available unlocked entry and locks it. If the transaction is blocked for more than 5000 milliseconds, then the method returns null.

```
Session session = ...;
ObjectMap map = session.getMap("xxx");
// this needs to be set somewhere to stop this loop
boolean timeToStop = false;

while(!timeToStop)
{
```

```

session.begin();
Object msgKey = map.getNextKey(5000);
if(msgKey == null)
{
    // current partition is exhausted, call it again in
    // a new transaction to move to next partition
    session.rollback();
    continue;
}
Message m = (Message)map.get(msgKey);
// now consume the message
...
// need to remove it
map.remove(msgKey);
session.commit();
}

```

Local mode versus client mode

If the application is using a local core, that is, it is not a client, then the mechanism works as described previously.

For client mode, if the Java virtual machine (JVM) is a client, then the client initially connects to a random partition primary. If no work exists in that partition, then the client moves to the next partition to look for work. The client either finds a partition with entries or loops around to the initial random partition. If the client loops around to the initial partition, then it returns a null value to the application. If the client finds a partition with a map that has entries, then it consumes entries from there until no entries are available for the timeout period. After the timeout passes, then null is returned. This action means that when null is returned and a partitioned map is used, then it you should start a new transaction and resume listening. The previous code sample fragment has this behavior.

Example

When you are running as a client and a key is returned, that transaction is now bound to the partition with the entry for that key. If you do not want to update any other maps during that transaction, then a problem does not exist. If you do want to update, then you can only update maps from the same partition as the map from which you got the key. The entry that is returned from the `getNextKey` method needs to give the application a way to discover relevant data in that partition. As an example, if you have two maps; one for events and another for jobs that the events impact. You define the two maps with the following entities:

Job.java

```

package tutorial.fifo;

import com.ibm.websphere.projector.annotations.Entity;
import com.ibm.websphere.projector.annotations.Id;

@Entity
public class Job
{
    @Id String jobId;

    int jobState;
}

```

JobEvent.java

```

package tutorial.fifo;

import com.ibm.websphere.projector.annotations.Entity;
import com.ibm.websphere.projector.annotations.Id;

```

```

import com.ibm.websphere.projector.annotations.OneToOne;

@Entity
public class JobEvent
{
    @Id String eventId;
    @OneToOne Job job;
}

```

The job has an ID and state, which is an integer. Suppose you want to increment the state whenever an event arrived. The events are stored in the JobEvent Map. Each entry has a reference to the job the event concerns. The code for the listener to do this looks like the following example:

JobEventListener.java

```

package tutorial.fifo;

import com.ibm.websphere.objectgrid.ObjectGridException;
import com.ibm.websphere.objectgrid.ObjectMap;
import com.ibm.websphere.objectgrid.Session;
import com.ibm.websphere.objectgrid.em.EntityManager;

public class JobEventListener
{
    boolean stopListening;

    public synchronized void stopListening()
    {
        stopListening = true;
    }

    synchronized boolean isStopped()
    {
        return stopListening;
    }

    public void processJobEvents(Session session)
        throws ObjectGridException
    {
        EntityManager em = session.getEntityManager();
        ObjectMap jobEvents = session.getMap("JobEvent");
        while(!isStopped())
        {
            em.getTransaction().begin();

            Object jobEventKey = jobEvents.getNextKey(5000);
            if(jobEventKey == null)
            {
                em.getTransaction().rollback();
                continue;
            }
            JobEvent event = (JobEvent)em.find(JobEvent.class, jobEventKey);
            // process the event, here we just increment the
            // job state
            event.job.jobState++;
            em.getTransaction().commit();
        }
    }
}

```

The listener is started on a thread by the application. The listener runs until the stopListening method is called. The processJobEvents method is run on the thread until the stopListening method is called. The loop blocks waiting for an eventKey from the JobEvent Map and then uses the EntityManager to access the event object, dereference to the job and increment the state.

The EntityManager API does not have a getNextKey method, but the ObjectMap does. So, the code uses the ObjectMap for JobEvent to get the key. If a map is used with entities then it does not store objects anymore. Instead, it stores Tuples; a Tuple object for the key and a Tuple object for the value. The EntityManager.find method accepts a Tuple for the key.

The code to create an event looks like the following example:

```
em.getTransaction().begin();
Job job = em.find(Job.class, "Job Key");
JobEvent event = new JobEvent();
event.id = Random.toString();
event.job = job;
em.persist(event); // insert it
em.getTransaction().commit();
```

You find the job for the event, construct an event, point it to the job, insert it in the JobEvent Map and commit the transaction.

Loaders and FIFO maps

If you want to back a map that is used as a FIFO queue with a Loader, then you might need to do some additional work. If the order of the entries in the map is not a concern, you have no extra work. If the order is important, then you need to add a sequence number to all of the inserted records when you are persisting the records to the backend. The preload mechanism should be written to insert the records on startup using this order.

Caching objects and their relationships (EntityManager API)

Most cache products use map-based APIs to store data as key-value pairs. The ObjectMap API and the dynamic cache in WebSphere Application Server, among others, use this approach. However, map-based APIs have limitations. The EntityManager API simplifies the interaction with the data grid by providing an easy way to declare and interact with a complex graph of related objects.

Map-based API limitations

If you are using a map-based API, such as the dynamic cache in WebSphere Application Server or the ObjectMap API, take the following limitations into consideration:

- Indexes and queries must use reflection to query fields and properties in cache objects.
- Custom data serialization is required to achieve performance for complex objects.
- It is difficult to work with graphs of objects. The application must store artificial references between objects and manually join the objects together.

Benefits of the EntityManager API

The EntityManager API uses the existing map-based infrastructure, but it converts entity objects to and from tuples before storing or reading them from the map. An entity object is transformed into a key tuple and a value tuple, which are then stored as key-value pairs. A tuple is an array of primitive attributes.

This set of APIs follows the Plain Old Java Object (POJO) style of programming that is adopted by most frameworks.

Defining an entity schema

An ObjectGrid can have any number of logical entity schemas. Entities are defined using annotated Java classes, XML, or a combination of both XML and Java classes. Defined entities are then registered with an eXtreme Scale server and bound to BackingMaps, indexes and other plug-ins.

When designing an entity schema, you must complete the following tasks:

1. Define the entities and their relationships.
2. Configure eXtreme Scale.
3. Register the entities.
4. Create entity-based applications that interact with the eXtreme Scale EntityManager APIs.

Entity schema configuration

An entity schema is a set of entities and the relationships between the entities. In an eXtreme Scale application with multiple partitions, the following restrictions and options apply to entity schemas:

- Each entity schema must have a single root defined. This is known as the schema root.
- All the entities for a given schema must be in the same map set, which means that all the entities that are reachable from a schema root with key or non-key relationships must be defined in the same map set as the schema root.
- Each entity can belong to only one entity schema.
- Each eXtreme Scale application can have multiple schemas.

Entities are registered with an ObjectGrid instance before it is initialized. Each defined entity must be uniquely named and is automatically bound to an ObjectGrid BackingMap of the same name. The initialization method varies depending on the configuration you are using:

Local eXtreme Scale configuration

If you are using a local ObjectGrid, you can programmatically configure the entity schema. In this mode, you can use the ObjectGrid.registerEntities methods to register annotated entity classes or an entity metadata descriptor file.

Distributed eXtreme Scale configuration

If you are using a distributed eXtreme Scale configuration, you must provide an entity metadata descriptor file with the entity schema.

For more details, see “EntityManager in a distributed environment” on page 71.

Entity requirements

Entity metadata is configured using Java class files, an entity descriptor XML file or both. At minimum, the entity descriptor XML is required to identify which eXtreme Scale BackingMaps are to be associated with entities. The persistent attributes of the entity and its relationships to other entities are described in either an annotated Java class (entity metadata class) or the entity descriptor XML file. The entity metadata class, when specified, is also used by the EntityManager API to interact with the data in the grid.

7.0.0.0 FIX 2+ An eXtreme Scale grid can be defined without providing any entity classes. This can be beneficial when the server and client are interacting directly with the tuple data stored in the underlying maps. Such entities are defined completely in the entity descriptor XML file and are referred to as classless entities.

Classless entities

Classless entities are useful when it is not possible to include application classes in the server or client classpath. Such entities are defined in the entity metadata descriptor XML file, where the class name of the entity is specified using a classless entity identifier in the form: @<entity identifier>. The @ symbol identifies the entity as classless and is used for mapping associations between entities. See the "Classless entity metadata" figure an example of an entity metadata descriptor XML file with two classless entities defined.

If an eXtreme Scale server or client does not have access to the classes, either can still use the EntityManager API using classless entities. Common use cases include the following:

- The eXtreme Scale container is hosted in a server that does not allow application classes in the classpath. In this case, the clients can still access the grid using the EntityManager API from a client, where the classes are allowed.
- The eXtreme Scale client does not require access to the entity classes because the client is either using a non-Java client, such as the eXtreme Scale REST data service or the client is accessing the tuple data in the grid using the ObjectMap API.

If the entity metadata is compatible between the client and server, entity metadata can be created using entity metadata classes, an XML file, or both.

For example, the "Programmatic entity class" in the following figure is compatible with the classless metadata code in the next section.

Programmatic entity class

```
@Entity
public class Employee {
    @Id long serialNumber;
    @Basic byte[] picture;
    @Version int ver;
    @ManyToOne(fetch=FetchType.EAGER, cascade=CascadeType.PERSIST)
    Department department;
}

@Entity
public static class Department {
    @Id int number;
    @Basic String name;
    @OneToMany(fetch=FetchType.LAZY, cascade=CascadeType.ALL, mappedBy="department")
    Collection<Employee> employees;
}
```

Classless fields, keys, and versions

As previously mentioned, classless entities are configured completely in the entity XML descriptor file. Class-based entities define their attributes using Java fields, properties and annotations. So classless entities need to define key and attribute structure in the entity XML descriptor with the <basic> and <id> tags.

Classless entity metadata

```
<?xml version="1.0" encoding="UTF-8"?>
<entity-mappings xmlns="http://ibm.com/ws/projector/config/emd"
```



```

xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/ws/projector/config/emd ./emd.xsd">

<entity class-name="@Employee" name="Employee">
  <attributes>
    <id name="serialNumber" type="long"/>
    <basic name="firstName" type="java.lang.String"/>
    <basic name="picture" type="[B"/>
    <version name="ver" type="int"/>
    <many-to-one
      name="department"
      target-entity="@Department"
      fetch="EAGER">
      <cascade><cascade-persist/></cascade>
    </many-to-one>
  </attributes>
</entity>

<entity class-name="@Department" name="Department" >
  <attributes>
    <id name="number" type="int"/>
    <basic name="name" type="java.lang.String"/>
    <version name="ver" type="int"/>
    <one-to-many
      name="employees"
      target-entity="@Employee"
      fetch="LAZY"
      mapped-by="department">
      <cascade><cascade-all/></cascade>
    </one-to-many>
  </attributes>
</entity>

```

Note that each entity above has an `<id>` element. A classless entity must have either one or more of an `<id>` element defined, or a single-valued association that represents the key for the entity. The fields of the entity are represented by `<basic>` elements. The `<id>`, `<version>`, and `<basic>` elements require a name and type in classless entities. See the following supported attribute types section for details on supported types.

Entity class requirements

Class-based entities are identified by associating various metadata with a Java class. The metadata can be specified using Java Platform, Standard Edition 5 annotations, an entity metadata descriptor file, or a combination of annotations and the descriptor file. Entity classes must meet the following criteria:

- The `@Entity` annotation is specified in the entity XML descriptor file.
- The class has a public or protected no-argument constructor.
- It must be a top-level class. Interfaces and enumerated types are not valid entity classes.
- Cannot use the final keyword.
- Cannot use inheritance.
- Must have a unique name and type for each ObjectGrid instance.

Entities all have a unique name and type. The name, if using annotations, is the simple (short) name of the class by default, but can be overridden using the name attribute of the `@Entity` annotation.

Persistent attributes

The persistent state of an entity is accessed by clients and the entity manager by using either fields (instance variables) or Enterprise JavaBeans-style property accessors. Each entity must define either field- or property-based access. Annotated entities are field-access if the class fields are annotated and are property-access if the getter method of the property is annotated. A mixture of field- and property-access is not allowed. If the type cannot be automatically determined, the **accessType** attribute on the `@Entity` annotation or equivalent XML can be used to identify the access type.

Persistent fields

Field-access entity instance variables are accessed directly from the entity manager and clients. Fields that are marked with the transient modifier or transient annotation are ignored. Persistent fields must not have final or static modifiers.

Persistent properties

Property-access entities must adhere to the JavaBeans signature conventions for read and write properties. Methods that do not follow JavaBeans conventions or have the Transient annotation on the getter method are ignored. For a property of type T, there must be a getter method `getProperty` which returns a value of type T and a void setter method `setProperty(T)`. For boolean types, the getter method can be expressed as `isProperty`, returning true or false. Persistent properties cannot have the static modifier.

Supported attribute types

The following persistent field and property types are supported:

- Java primitive types including wrappers
- `java.lang.String`
- `java.math.BigInteger`
- `java.math.BigDecimal`
- `java.util.Date`
- `java.util.Calendar`
- `java.sql.Date`
- `java.sql.Time`
- `java.sql.Timestamp`
- `byte[]`
- `java.lang.Byte[]`
- `char[]`
- `java.lang.Character[]`
- `enum`

User serializable attribute types are supported but have performance, query and change-detection limitations. Persistent data that cannot be proxied, such as arrays and user serializable objects, must be reassigned to the entity if altered.

Serializable attributes are represented in the entity descriptor XML file using the class name of the object. If the object is an array, the data type is represented using the Java internal form. For example, if an attribute data type is `java.lang.Byte[][]`, the string representation is `[[Ljava.lang.Byte;`

User serializable types should adhere to the following best practices:

- Implement high performance serialization methods. Implement the `java.lang.Cloneable` interface and `public clone` method.
- Implement the `java.io.Externalizable` interface.
- Implement `equals` and `hashCode`

Entity associations

Bi-directional and uni-directional entity associations, or relationships between entities can be defined as one-to-one, many-to-one, one-to-many and many-to-many. The entity manager automatically resolves the entity relationships to the appropriate key references when storing the entities.

The eXtreme Scale grid is a data cache and does not enforce referential integrity like a database. Although relationships allow cascading persist and remove operations for child entities, it does not detect or enforce broken links to objects. When removing a child object, the reference to that object must be removed from the parent.

If you define a bi-directional association between two entities, you must identify the owner of the relationship. In a to-many association, the many side of the relationship is always the owning side. If ownership cannot be determined automatically, then the **mappedBy** attribute of the annotation, or XML equivalent, must be specified. The **mappedBy** attribute identifies the field in the target entity that is the owner of the relationship. This attribute also helps identify the related fields when there are multiple attributes of the same type and cardinality.

Single-valued associations

One-to-one and many-to-one associations are denoted using the `@OneToOne` and `@ManyToOne` annotations or equivalent XML attributes. The target entity type is determined by the attribute type. The following example defines a uni-directional association between `Person` and `Address`. The `Customer` entity has a reference to one `Address` entity. In this case, the association could also be many-to-one since there is no inverse relationship.

```
@Entity
public class Customer {
    @Id id;
    @OneToOne Address homeAddress;
}

@Entity
public class Address{
    @Id id
    @Basic String city;
}
```

To specify a bi-directional relationship between the `Customer` and `Address` classes, add a reference to the `Customer` class from the `Address` class and add the appropriate annotation to mark the inverse side of the relationship. Because this association is one-to-one, you have to specify an owner of the relationship using the `mappedBy` attribute on the `@OneToOne` annotation.

```

@Entity
public class Address{
    @Id id
    @Basic String city;
    @OneToOne(mappedBy="homeAddress") Customer customer;
}

```

Collection-valued associations

One-to-many and many-to-many associations are denoted using the `@OneToMany` and `@ManyToMany` annotations or equivalent XML attributes. All many relationships are represented using the types: `java.util.Collection`, `java.util.List` or `java.util.Set`. The target entity type is determined by the generic type of the `Collection`, `List` or `Set` or explicitly using the **targetEntity** attribute on the `@OneToMany` or `@ManyToMany` annotation (or XML equivalent).

In the previous example, it is not practical to have one address object per customer because many customers might share an address or might have multiple addresses. This situation is better solved using a many association:

```

@Entity
public class Customer {
    @Id id;
    @ManyToOne Address homeAddress;
    @ManyToOne Address workAddress;
}

@Entity
public class Address{
    @Id id
    @Basic String city;
    @OneToMany(mappedBy="homeAddress") Collection<Customer> homeCustomers;

    @OneToMany(mappedBy="workAddress", targetEntity=Customer.class)
    Collection workCustomers;
}

```

In this example, two different relationships exist between the same entities: a Home and Work address relationship. A non-generic `Collection` is used for the **workCustomers** attribute to demonstrate how to use the **targetEntity** attribute when generics are not available.

Classless associations

Classless entity associations are defined in the entity metadata descriptor XML file similar to how class-based associations are defined. The only difference is that instead of the target entity pointing to an actual class, it points to the classless entity identifier used for the class name of the entity.

An example follows:

```

<many-to-one name="department" target-entity="@Department" fetch="EAGER">
    <cascade><cascade-all/></cascade>
</many-to-one>
<one-to-many name="employees" target-entity="@Employee" fetch="LAZY">
    <cascade><cascade-all/></cascade>
</one-to-many>

```

Primary keys

All entities must have a primary key, which can be a simple (single attribute) or composite (multiple attribute) key. The key attributes are denoted using the `Id` annotation or defined in the entity XML descriptor file. Key attributes have the following requirements:

- The value of a primary key cannot change.
- A primary key attribute should be one of the following types: Java primitive type and wrappers, `java.lang.String`, `java.util.Date` or `java.sql.Date`.
- A primary key can contain any number of single-valued associations. The target entity of the primary key association must not have an inverse association directly or indirectly to the source entity.

Composite primary keys can optionally define a primary key class. An entity is associated with a primary key class using the `@IdClass` annotation or the entity XML descriptor file. An `@IdClass` annotation is useful in conjunction with the `EntityManager.find` method.

Primary key classes have the following requirements:

- It should be public with a no-argument constructor.
- The access type of the primary key class is determined by the entity that declares the primary key class.
- If property-access, the properties of the primary key class must be public or protected.
- The primary key fields or properties must match the key attribute names and types defined in the referencing entity.
- Primary key classes must implement the `equals` and `hashCode` methods.

An example follows:

```
@Entity
@IdClass(CustomerKey.class)
public class Customer {
    @Id @ManyToOne Zone zone;
    @Id int custId;
    String name;
    ...
}

@Entity
public class Zone{
    @Id String zoneCode;
    String name;
}

public class CustomerKey {
    Zone zone;
    int custId;

    public int hashCode() {...}
    public boolean equals(Object o) {...}
}
```

Classless primary keys

Classless entities are required to either have at least one `<id>` element or an association in the XML file with the attribute `id=true`. An example of both would look like the following:

```

<id name="serialNumber" type="int"/>
<many-to-one name="department" target-entity="@Department" id="true">
<cascade><cascade-all/></cascade>
</many-to-one>

```

Remember:

The <id-class> XML tag is not supported for classless entities.

Entity proxies and field interception

Entity classes and mutable supported attribute types are extended by proxy classes for property-access entities and bytecode-enhanced for Java Development Kit (JDK) 5 field-access entities. All access to the entity, even by internal business methods and the equals methods, must use the appropriate field or property access methods.

Proxies and field interceptors are used to allow the entity manager to track the state of the entity, determine if the entity has changed, and improve performance. Field interceptors are only available on Java SE 5 platforms when the entity instrumentation agent is configured.

Attention: When using property-access entities, the equals method should use the instanceof operator for comparing the current instance to the input object. All introspection of the target object should be through the properties of the object, not the fields themselves, because the object instance will be the proxy.

emd.xsd file

Use the entity metadata XML schema definition to create a descriptor XML file and define an entity schema for WebSphere eXtreme Scale.

See the the information about the entity metadata descriptor file in the *Administration Guide* for the descriptions of each element and attribute of the emd.xsd file.

emd.xsd file

```

<?xml version="1.0" encoding="UTF-8"?>
<xsd:schema xmlns:emd="http://ibm.com/ws/projector/config/emd"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  targetNamespace="http://ibm.com/ws/projector/config/emd"
  elementFormDefault="qualified" attributeFormDefault="unqualified"
  version="1.0">

  <!-- ***** -->
  <xsd:element name="entity-mappings">
    <xsd:complexType>
      <xsd:sequence>
        <xsd:element name="description" type="xsd:string" minOccurs="0" />
        <xsd:element name="entity" type="emd:entity" minOccurs="1" maxOccurs="unbounded" />
      </xsd:sequence>
    </xsd:complexType>
    <xsd:unique name="uniqueEntityClassName">
      <xsd:selector xpath="emd:entity" />
      <xsd:field xpath="class-name" />
    </xsd:unique>
  </xsd:element>

  <!-- ***** -->
  <xsd:complexType name="entity">
    <xsd:sequence>
      <xsd:element name="description" type="xsd:string" minOccurs="0" />
      <xsd:element name="id-class" type="emd:id-class" minOccurs="0" />
      <xsd:element name="attributes" type="emd:attributes" minOccurs="0" />
      <xsd:element name="entity-listeners" type="emd:entity-listeners" minOccurs="0" />
      <xsd:element name="pre-persist" type="emd:pre-persist" minOccurs="0" />
      <xsd:element name="post-persist" type="emd:post-persist" minOccurs="0" />
      <xsd:element name="pre-remove" type="emd:pre-remove" minOccurs="0" />
      <xsd:element name="post-remove" type="emd:post-remove" minOccurs="0" />
      <xsd:element name="pre-invalidate" type="emd:pre-invalidate" minOccurs="0" />
      <xsd:element name="post-invalidate" type="emd:post-invalidate" minOccurs="0" />
      <xsd:element name="pre-update" type="emd:pre-update" minOccurs="0" />
    </xsd:sequence>
  </xsd:complexType>

```

```

        <xsd:element name="post-update" type="emd:post-update" minOccurs="0" />
        <xsd:element name="post-load" type="emd:post-load" minOccurs="0" />
    </xsd:sequence>
    <xsd:attribute name="name" type="xsd:string" use="required" />
    <xsd:attribute name="class-name" type="xsd:string" use="required" />
    <xsd:attribute name="access" type="emd:access-type" />
    <xsd:attribute name="schemaRoot" type="xsd:boolean" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="attributes">
    <xsd:sequence>
        <xsd:choice>
            <xsd:element name="id" type="emd:id" minOccurs="0" maxOccurs="unbounded" />
        </xsd:choice>
        <xsd:element name="basic" type="emd:basic" minOccurs="0" maxOccurs="unbounded" />
        <xsd:element name="version" type="emd:version" minOccurs="0" maxOccurs="unbounded"/>
        <xsd:element name="many-to-one" type="emd:many-to-one" minOccurs="0" maxOccurs="unbounded" />
        <xsd:element name="one-to-many" type="emd:one-to-many" minOccurs="0" maxOccurs="unbounded" />
        <xsd:element name="one-to-one" type="emd:one-to-one" minOccurs="0" maxOccurs="unbounded" />
        <xsd:element name="many-to-many" type="emd:many-to-many" minOccurs="0" maxOccurs="unbounded" />
        <xsd:element name="transient" type="emd:transient" minOccurs="0" maxOccurs="unbounded" />
    </xsd:sequence>
</xsd:complexType>

<!-- ***** -->
<xsd:simpleType name="access-type">
    <xsd:restriction base="xsd:token">
        <xsd:enumeration value="PROPERTY" />
        <xsd:enumeration value="FIELD" />
    </xsd:restriction>
</xsd:simpleType>

<!-- ***** -->
<xsd:complexType name="id-class">
    <xsd:attribute name="class-name" type="xsd:string" use="required" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="id">
    <xsd:attribute name="name" type="xsd:string" use="required" />
    <xsd:attribute name="type" type="xsd:string" />
    <xsd:attribute name="alias" type="xsd:string" use="optional" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="transient">
    <xsd:attribute name="name" type="xsd:string" use="required" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="basic">
    <xsd:attribute name="name" type="xsd:string" use="required" />
    <xsd:attribute name="alias" type="xsd:string" />
    <xsd:attribute name="type" type="xsd:string" />
    <xsd:attribute name="fetch" type="emd:fetch-type" />
</xsd:complexType>

<!-- ***** -->
<xsd:simpleType name="fetch-type">
    <xsd:restriction base="xsd:token">
        <xsd:enumeration value="LAZY" />
        <xsd:enumeration value="EAGER" />
    </xsd:restriction>
</xsd:simpleType>

<!-- ***** -->
<xsd:complexType name="many-to-one">
    <xsd:sequence>
        <xsd:element name="cascade" type="emd:cascade-type" minOccurs="0" />
    </xsd:sequence>
    <xsd:attribute name="name" type="xsd:string" use="required" />
    <xsd:attribute name="alias" type="xsd:string" />
    <xsd:attribute name="target-entity" type="xsd:string" />
    <xsd:attribute name="fetch" type="emd:fetch-type" />
    <xsd:attribute name="id" type="xsd:boolean" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="one-to-one">
    <xsd:sequence>
        <xsd:element name="cascade" type="emd:cascade-type" minOccurs="0" />
    </xsd:sequence>
    <xsd:attribute name="name" type="xsd:string" use="required" />
    <xsd:attribute name="alias" type="xsd:string" />
    <xsd:attribute name="target-entity" type="xsd:string" />
    <xsd:attribute name="fetch" type="emd:fetch-type" />
    <xsd:attribute name="mapped-by" type="xsd:string" />
    <xsd:attribute name="id" type="xsd:boolean" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="one-to-many">

```

```

<xsd:sequence>
  <xsd:element name="order-by" type="emd:order-by" minOccurs="0" />
  <xsd:element name="cascade" type="emd:cascade-type" minOccurs="0" />
</xsd:sequence>
<xsd:attribute name="name" type="xsd:string" use="required" />
<xsd:attribute name="alias" type="xsd:string" />
<xsd:attribute name="target-entity" type="xsd:string" />
<xsd:attribute name="fetch" type="emd:fetch-type" />
<xsd:attribute name="mapped-by" type="xsd:string" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="many-to-many">
  <xsd:sequence>
    <xsd:element name="order-by" type="emd:order-by" minOccurs="0" />
    <xsd:element name="cascade" type="emd:cascade-type" minOccurs="0" />
  </xsd:sequence>
  <xsd:attribute name="name" type="xsd:string" use="required" />
  <xsd:attribute name="alias" type="xsd:string" />
  <xsd:attribute name="target-entity" type="xsd:string" />
  <xsd:attribute name="fetch" type="emd:fetch-type" />
  <xsd:attribute name="mapped-by" type="xsd:string" />
</xsd:complexType>

<!-- ***** -->
<xsd:simpleType name="order-by">
  <xsd:restriction base="xsd:string" />
</xsd:simpleType>

<!-- ***** -->
<xsd:complexType name="cascade-type">
  <xsd:sequence>
    <xsd:element name="cascade-all" type="emd:emptyType" minOccurs="0" />
    <xsd:element name="cascade-persist" type="emd:emptyType" minOccurs="0" />
    <xsd:element name="cascade-remove" type="emd:emptyType" minOccurs="0" />
    <xsd:element name="cascade-invalidate" type="emd:emptyType" minOccurs="0" />
    <xsd:element name="cascade-merge" type="emd:emptyType" minOccurs="0" />
    <xsd:element name="cascade-refresh" type="emd:emptyType" minOccurs="0" />
  </xsd:sequence>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="emptyType" />

<!-- ***** -->
<xsd:complexType name="version">
  <xsd:attribute name="name" type="xsd:string" use="required"/>
  <xsd:attribute name="alias" type="xsd:string" />
  <xsd:attribute name="type" type="xsd:string" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="entity-listeners">
  <xsd:sequence>
    <xsd:element name="entity-listener" type="emd:entity-listener" minOccurs="0" maxOccurs="unbounded" />
  </xsd:sequence>
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="entity-listener">
  <xsd:sequence>
    <xsd:element name="pre-persist" type="emd:pre-persist" minOccurs="0" />
    <xsd:element name="post-persist" type="emd:post-persist" minOccurs="0" />
    <xsd:element name="pre-remove" type="emd:pre-remove" minOccurs="0" />
    <xsd:element name="post-remove" type="emd:post-remove" minOccurs="0" />
    <xsd:element name="pre-invalidate" type="emd:pre-invalidate" minOccurs="0" />
    <xsd:element name="post-invalidate" type="emd:post-invalidate" minOccurs="0" />
    <xsd:element name="pre-update" type="emd:pre-update" minOccurs="0" />
    <xsd:element name="post-update" type="emd:post-update" minOccurs="0" />
    <xsd:element name="post-load" type="emd:post-load" minOccurs="0" />
  </xsd:sequence>
  <xsd:attribute name="class-name" type="xsd:string" use="required" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="pre-persist">
  <xsd:attribute name="method-name" type="xsd:string" use="required" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="post-persist">
  <xsd:attribute name="method-name" type="xsd:string" use="required" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="pre-remove">
  <xsd:attribute name="method-name" type="xsd:string" use="required" />
</xsd:complexType>

<!-- ***** -->

```



```

<xsd:complexType name="post-remove">
  <xsd:attribute name="method-name" type="xsd:string" use="required" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="pre-invalidate">
  <xsd:attribute name="method-name" type="xsd:string" use="required" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="post-invalidate">
  <xsd:attribute name="method-name" type="xsd:string" use="required" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="pre-update">
  <xsd:attribute name="method-name" type="xsd:string" use="required" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="post-update">
  <xsd:attribute name="method-name" type="xsd:string" use="required" />
</xsd:complexType>

<!-- ***** -->
<xsd:complexType name="post-load">
  <xsd:attribute name="method-name" type="xsd:string" use="required" />
</xsd:complexType>

</xsd:schema>

```

EntityManager in a distributed environment

You can use EntityManager with a local ObjectGrid or in a distributed eXtreme Scale environment. The main difference is how you connect to this remote environment. After you establish a connection, there is no difference between using a Session object or using the EntityManager API.

Required configuration files

The following XML configuration files are required:

- ObjectGrid descriptor XML file
- Entity descriptor XML file
- Deployment or data grid descriptor XML file

These files specify which entities and BackingMaps a server will host.

The entity metadata descriptor file contains a description of the entities that are used. At minimum, you must specify the entity class and name. If you are running in a Java Platform, Standard Edition 5 environment, eXtreme Scale automatically reads the entity class and its annotations. You can define additional XML attributes if the entity class has no annotations or if you are required to override the class attributes. If you are registering the entities classless, provide all of entity information in the XML file only.

You can use the following XML configuration snippet to define a data grid with entities. In this snippet, the server creates an ObjectGrid with the name bookstore and an associated backing map with the name order. Note that the objectgrid.xml file snippet refers to the entity.xml file. In this case, the entity.xml file contains one entity, the Order entity.

objectgrid.xml

```

<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">

  <objectGrids>
    <objectGrid name="bookstore" entityMetadataXMLFile="entity.xml">
      <backingMap name="Order"/>
    </objectGrid>
  </objectGrids>

```

```

        </objectGrid>
    </objectGrids>

</objectGridConfig>

```

This objectgrid.xml file refers to the entity.xml with the **entityMetadataXMLFile** attribute. The location of this file is relative to the location of the objectgrid.xml file. An example of the entity.xml file follows:

```

entity.xml
<entity-mappings xmlns="http://ibm.com/ws/projector/config/emd"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/projector/config/emd ./emd.xsd">
  <entity class-name="com.ibm.websphere.tutorials.objectgrid.em.
    distributed.step1.Order" name="Order"/>
</entity-mappings>

```

This example assumes the Order class would have the orderNumber and desc fields annotated similarly.

An equivalent classless entity.xml file would be as follows

```

classless entity.xml
<entity-mappings xmlns="http://ibm.com/ws/projector/config/emd"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/projector/config/emd ./emd.xsd">
  <entity class-name="@Order " name="Order">
    <description>"Entity named: Order"</description>
    <attributes>
      <id name="orderNumber" type="int"/>
      <basic name="desc" type="java.lang.String"/>
    </attributes>
  </entity>
</entity-mappings>

```

For information about starting an eXtreme Scale server, see *Starting WebSphere eXtreme Scale server processes* in the *Administration Guide*, which uses both the deployment.xml and objectgrid.xml files to start the catalog server.

Connecting to a distributed eXtreme Scale server

The following code enables the connect mechanism for a client and server on the same computer:

```

String catalogEndpoints="localhost:2809";
URL clientOverrideURL= new URL("file:etc/emtutorial/distributed/step1/objectgrid.xml");
ClientClusterContext clusterCtx = ogMgr.connect(catalogEndpoints, null, clientOverrideURL);
ObjectGrid objectGrid=ogMgr.getObjectGrid(clusterCtx, "bookstore");

```

In the preceding code snippet, note the reference to the remote eXtreme Scale server. After you establish a connection , you can invoke EntityManager API methods such as persist, update, remove and find.

Attention: When you are using entities, pass the client override ObjectGrid descriptor XML file to the connect method. If a null value is passed to the clientOverrideURL property and the client has a different directory structure than the server, then the client might fail to locate the ObjectGrid or entity descriptor XML files. At minimum, the ObjectGrid and entity XML files for the server can be copied to the client.

Previously, using entities on an ObjectGrid client required you to make the ObjectGrid XML and entity XML available to the client in one of the following two ways:

1. Pass an overriding ObjectGrid XML to the ObjectGridManager.connect(String catalogServerAddresses, ClientSecurityConfiguration securityProps, URL overRideObjectGridXml) method.

```
String catalogEndpoints="myHost:2809";
URL clientOverrideURL= new URL("file:etc/emtutorial/distributed/step1/objectgrid.xml");
ClientClusterContext clusterCtx = ogMgr.connect(catalogEndpoints, null, clientOverrideURL);
ObjectGrid objectGrid=ogMgr.getObjectGrid(clusterCtx, "bookstore");
```

2. Pass null for the override file and ensure that the ObjectGrid XML and referenced entity XML are available to the client on the same path as on the server.

```
String catalogEndpoints="myHost:2809";
ClientClusterContext clusterCtx = ogMgr.connect(catalogEndpoints, null, null);
ObjectGrid objectGrid=ogMgr.getObjectGrid(clusterCtx, "bookstore");
```

7.0.0.0 FIX 2+ The XML files were required regardless of whether or not you wanted to use subset entities on the client side. These files are no longer required to use the entities as defined by the server. Instead, pass null as the `overRideObjectGridXml` parameter as in option 2 of the previous section. If the XML file is not found on the same path set on the server, the client will use the entity configuration on the server.

However, if you use subset entities on the client, you must provide an overriding ObjectGrid XML as in option 1.

Client and server side schema

The server-side schema defines the type of data stored in the maps on a server. The client-side schema is a mapping to application objects from the schema on the server. For example, you might have the following server-side schema:

```
@Entity
class ServerPerson
{
    @Id String ssn;
    String firstName;
    String surname;
    int age;
    int salary;
}
```

A client can have an object annotated as in the following example:

```
@Entity(name="ServerPerson")
class ClientPerson
{
    @Id @Basic(alias="ssn") String socialSecurityNumber;
    String surname;
}
```

This client then takes a server-side entity and projects the subset of the entity into the client object. This projection leads to bandwidth and memory savings on a client because the client has only the information it needs instead of all of the information that is in the server side entity. Different applications can use their own objects instead of forcing all applications to share a set of classes for data access.

The client-side entity descriptor XML file is required in the following cases: if the server is running with class-based entities while the client side is running classless; or if the server is classless and the client uses class-based entities. A classless client mode allows the client to still run entity queries without having access to the physical classes. Assuming the server has registered the `ServerPerson` entity above, the client would override the data grid with an `entity.xml` file such as follows:

```
<entity-mappings xmlns="http://ibm.com/ws/projector/config/emd"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/ws/projector/config/emd ./emd.xsd">
<entity class-name="@ServerPerson" name="Order">
<description>"Entity named: Order"</description>
<attributes>
<id name="socialSecurityNumber" type="java.lang.String"/>
```

```
        <basic name="surname" type="java.lang.String"/>
      </attributes>
    </entity>
  </entity-mappings>
```

This file achieves an equivalent subset entity on the client, without requiring the client to provide the actual annotated class. If the server is classless, and the client is not classless, then the client provides an overriding entity descriptor XML file. This entity descriptor XML file contains an override to the class file reference.

Referencing the schema

If your application is running in Java SE 5, then the application can be added to the objects by using annotations. The EntityManager can read the schema from the annotations on those objects. The application provides the eXtreme Scale run time with references to these objects using the `entity.xml` file, which is referenced from the `objectgrid.xml` file. The `entity.xml` file lists all the entities, each of which is associated with either a class or a schema. If a proper class name is specified, then the application attempts to read the Java SE 5 annotations from those classes to determine the schema. If you do not annotate the class file or specify a classless identifier as the class name, then the schema is taken from the XML file. The XML file is used to specify all the attributes, keys, and relationships for each entity.

A local data grid does not need XML files. The program can obtain an ObjectGrid reference and invoke the `ObjectGrid.registerEntities` method to specify a list of Java SE 5 annotated classes or an XML file.

The run time uses the XML file or a list of annotated classes to find entity names, attribute names and types, key fields and types, and relationships between entities. If eXtreme Scale is running on a server or in stand-alone mode, then it automatically makes a map named after each entity. These maps can be customized further using the `objectgrid.xml` file or APIs set either by the application or injection frameworks such as Spring.

Entity metadata descriptor file

See “`emd.xsd` file” on page 68 for more information about the metadata descriptor file.

Interacting with EntityManager

Applications typically first obtain an ObjectGrid reference, and then a Session from that reference for each thread. Sessions cannot be shared between threads. An extra method on Session, the `getEntityManager` method, is available. This method returns a reference to an entity manager to use for this thread. The EntityManager interface can replace the Session and ObjectMap interfaces for all applications. You can use these EntityManager APIs if the client has access to the defined entity classes.

Obtaining an EntityManager instance from a session

The `getEntityManager` method is available on a Session object. The following code example illustrates how to create a local ObjectGrid instance and access the EntityManager. See the EntityManager interface in the API documentation for details about all the supported methods.

```
ObjectGrid og =
ObjectGridManagerFactory.getObjectGridManager().createObjectGrid("intro-grid");
Session s = og.getSession();
EntityManager em = s.getEntityManager();
```

A one-to-one relationship exists between the Session object and EntityManager object. You can use the EntityManager object more than once.

Persisting an entity

Persisting an entity means saving the state of a new entity in an ObjectGrid cache. After the persist method is called, the entity is in the managed state. Persist is a transactional operation, and the new entity is stored in the ObjectGrid cache after the transaction commits.

Every entity has a corresponding BackingMap in which the tuples are stored. The BackingMap has the same name as the entity, and is created when the class is registered. The following code example demonstrates how to create an Order object by using the persist operation.

```
Order order = new Order(123);
em.persist(order);
order.setX();
...
```

The Order object is created with the key 123, and the object is passed to the persist method. You can continue to modify the state of the object before you commit the transaction.

Important: The preceding example does not include any required transactional boundaries, such as begin and commit. See the the entity manager tutorial in the *Product Overview* for more information.

Finding an entity

You can locate the entity in the ObjectGrid cache with the find method by providing a key after the entity is stored in the cache. This method does not require any transactional boundary, which is useful for read-only semantics. The following example illustrates that only one line of code is needed to locate the entity.

```
Order foundOrder = (Order)em.find(Order.class, new Integer(123));
```

Removing an entity

The remove method, like the persist method, is a transactional operation. The following example shows the transactional boundary by calling the begin and commit methods.

```
em.getTransaction().begin();
Order foundOrder = (Order)em.find(Order.class, new Integer(123));
em.remove(foundOrder );
em.getTransaction().commit();
```

The entity must first be managed before it can be removed, which you can accomplish by calling the find method within the transactional boundary. Then call the remove method on the EntityManager interface.

Invalidating an entity

The invalidate method behaves much like the remove method, but does not invoke any Loader plug-ins. Use this method to remove entities from the ObjectGrid, but to preserve them in the backend data store.

```
em.getTransaction().begin();
Order foundOrder = (Order)em.find(Order.class, new Integer(123));
em.invalidate(foundOrder);
em.getTransaction().commit();
```

The entity must first be managed before it can be invalidated, which you can accomplish by calling the find method within the transactional boundary. After you call the find method, you can call the invalidate method on the EntityManager interface.

Updating an entity

The update method is also a transactional operation. The entity must be managed before any updates can be applied.

```
em.getTransaction().begin();
Order foundOrder = (Order)em.find(Order.class, new Integer(123));
foundOrder.date = new Date(); // update the date of the order
em.getTransaction().commit();
```

In the preceding example, the persist method is not called after the entity is updated. The entity is updated in the ObjectGrid cache when the transaction is committed.

Queries and query queues

With the flexible query engine, you can retrieve entities by using EntityManager API. Create SELECT type queries over an entity or Object-based schema by using the ObjectGrid query language. Query interface explains in detail how you can run the queries by using the EntityManager API. See the Query API for more information about using queries.

An entity QueryQueue is a queue-like data structure associated with an entity query. It selects all the entities that match the WHERE condition on the query filter and puts the result entities in a queue. Clients can then iteratively retrieve entities from this queue. See “Entity query queues” on page 87 for more information.

Entity listeners and callback methods

Applications can be notified when the state of an entity transitions from state to state. Two callback mechanisms exist for state change events: life cycle callback methods that are defined on an entity class and are invoked whenever the entity state changes, and entity listeners, which are more general because the entity listener can be registered on several entities.

Life cycle of an entity instance

An entity instance has the following states:

- **New:** A newly created entity instance that does not exist in the eXtreme Scale cache.
- **Managed:** The entity instance exists in the eXtreme Scale cache and is retrieved or persisted using the entity manager. An entity must be associated with an active transaction to be in the managed state.

- **Detached:** The entity instance exists in the eXtreme Scale cache, but is no longer associated with an active transaction.
- **Removed:** The entity instance is removed, or is scheduled to be removed, from the eXtreme Scale cache when the transaction is flushed or committed.
- **Invalidated:** The entity instance is invalidated, or is scheduled to be invalidated, from the eXtreme Scale cache when the transaction is flushed or committed.

When entities change from state to state, you can invoke life-cycle, call-back methods.

The following sections further describe the meanings of New, Managed, Detached, Removed and Invalidated states as the states apply to an entity.

Entity life cycle callback methods

Entity life cycle callback methods can be defined on the entity class and are invoked when the entity state changes. These methods are useful for validating entity fields and updating transient state that is not usually persisted with the entity. Entity life cycle callback methods can also be defined on classes that are not using entities. Such classes are entity listener classes, which can be associated with multiple entity types. life cycle callback methods can be defined using both metadata annotations and a entity metadata XML descriptor file:

- **Annotations:** life cycle callback methods can be denoted using the PrePersist, PostPersist, PreRemove, PostRemove, PreUpdate, PostUpdate, and PostLoad annotations in an entity class.
- **Entity XML descriptor :** life cycle callback methods can be described using XML when annotations are not available.

Entity listeners

An entity listener class is a class that does not use entities that defines one or more entity life cycle callback methods. Entity listeners are useful for general purpose auditing or logging applications. Entity listeners can be defined using both metadata annotations and a entity metadata XML descriptor file:

- **Annotation:** The EntityListeners annotation can be used to denote one or more entity listener classes on an entity class. If multiple entity listeners are defined, the order in which they are invoked is determined by the order in which they are specified in the EntityListeners annotation.
- **Entity XML descriptor:** The XML descriptor can be used as an alternative to specify the invocation order of entity listeners or to override the order that is specified in metadata annotations.

Callback method requirements

Any subset or combination of annotations can be specified on an entity class or a listener class. A single class cannot have more than one life cycle callback method for the same life cycle event. However, the same method can be used for multiple callback events. The entity listener class must have a public no-arg constructor. Entity listeners are stateless. The life cycle of an entity listener is unspecified. eXtreme Scale does not support entity inheritance, so callback methods can only be defined in the entity class, but not in the superclass.

Callback method signature

Entity life cycle callback methods can be defined on an entity listener class, directly on an entity class, or both. Entity life cycle callback methods can be defined using both metadata annotations and the entity XML descriptor. The annotations used for callback methods on the entity class and on the entity listener class are the same. The signatures of the callback methods are different when defined on an entity class versus an entity listener class. Callback methods defined on an entity class or mapped superclass have the following signature:

```
void <METHOD>()
```

Callback methods that are defined on an entity listener class have the following signature:

```
void <METHOD>(Object)
```

The Object argument is the entity instance for which the callback method is invoked. The Object argument can be declared as a `java.lang.Object` object or the actual entity type.

Callback methods can have public, private, protected, or package level access, but must not be static or final.

The following annotations are defined to designate life cycle event callback methods of the corresponding types:

- `com.ibm.websphere.projector.annotations.PrePersist`
- `com.ibm.websphere.projector.annotations.PostPersist`
- `com.ibm.websphere.projector.annotations.PreRemove`
- `com.ibm.websphere.projector.annotations.PostRemove`
- `com.ibm.websphere.projector.annotations.PreUpdate`
- `com.ibm.websphere.projector.annotations.PostUpdate`
- `com.ibm.websphere.projector.annotations.PostLoad`

See the API Documentation for more details. Each annotation has an equivalent XML attribute defined in the entity metadata XML descriptor file.

Life cycle callback method semantics

Each of the different life cycle callback methods has a different purpose and is called in different phases of the entity life cycle:

PrePersist

Invoked for an entity before the entity has been persisted to the store, which includes entities that have been persisted due to a cascading operation. This method is invoked on the thread of the `EntityManager.persist` operation.

PostPersist

Invoked for an entity after the entity has been persisted to the store, which includes entities that have been persisted due to a cascading operation. This method is invoked on the thread of the `EntityManager.persist` operation. It is called after the `EntityManager.flush` or `EntityManager.commit` is called.

PreRemove

Invoked for an entity before the entity has been removed, which includes

entities that have been removed due to a cascading operation. This method is invoked on the thread of the EntityManager.remove operation.

PostRemove

Invoked for an entity after the entity has been removed, which includes entities that have been removed due to a cascading operation. This method is invoked on the thread of the EntityManager.remove operation. It is called after the EntityManager.flush or EntityManager.commit is called.

PreUpdate

Invoked for an entity before the entity has been updated to the store. This method is invoked on the thread of the transaction flush or commit operation.

PostUpdate

Invoked for an entity after the entity has been updated to the store. This method is invoked on the thread of the transaction flush or commit operation.

PostLoad

Invoked for an entity after the entity has been loaded from the store which includes any entities that are loaded through an association. This method is invoked on the thread of the loading operation, such as EntityManager.find or a query.

Duplicate life cycle callback methods

If multiple callback methods are defined for an entity life cycle event, the ordering of the invocation of these methods is as follows:

1. **life cycle callback methods defined in the entity listeners:** The life cycle callback methods that are defined on the entity listener classes for an entity class are invoked in the same order as the specification of the entity listener classes in the EntityListeners annotation or the XML descriptor.
2. **Listener super class:** Callback methods defined in the super class of the entity listener are invoked before the children.
3. **Entity life cycle methods:** WebSphere eXtreme Scale does not support entity inheritance, so the entity life cycle methods can only be defined in the entity class.

Exceptions

Life cycle callback methods might result in run time exceptions. If a life cycle callback method results in a run time exception within a transaction, the transaction is rolled back. No further life cycle callback methods are invoked after a runtime exception results.

Entity listener examples

You can write EntityListeners based on your requirements. Several example scripts follow.

EntityListeners example using annotations

The following example shows the life-cycle callback method invocations and order of the invocations. Assume an entity class Employee and two entity listeners exist: EmployeeListener and EmployeeListener2.

```

@Entity
@EntityListeners(EmployeeListener.class, EmployeeListener2.class)
public class Employee {
    @PrePersist
    public void checkEmployeeID() {
        ....
    }
}

public class EmployeeListener {
    @PrePersist
    public void onEmployeePrePersist(Employee e) {
        ....
    }
}

public class PersonListener {
    @PrePersist
    public void onPersonPrePersist(Object person) {
        ....
    }
}

public class EmployeeListener2 {
    @PrePersist
    public void onEmployeePrePersist2(Object employee) {
        ....
    }
}

```

If a PrePersist event occurs on an Employee instance, the following methods are called in order:

1. onEmployeePrePersist method
2. onPersonPrePersist method
3. onEmployeePrePersist2 method
4. checkEmployeeID method

Entity listeners example using XML

The following example shows how to set an entity listener on an entity using the entity descriptor XML file:

```

<entity
  class-name="com.ibm.websphere.objectgrid.sample.Employee"
  name="Employee" access="FIELD">
  <attributes>
    <id name="id" />
    <basic name="value" />
  </attributes>
  <entity-listeners>
    <entity-listener
      class-name="com.ibm.websphere.objectgrid.sample.EmployeeListener">
      <pre-persist method-name="onListenerPrePersist" />
      <post-persist method-name="onListenerPostPersist" />
    </entity-listener>
  </entity-listeners>
  <pre-persist method-name="checkEmployeeID" />
</entity>

```

The entity Employee is configured with a com.ibm.websphere.objectgrid.sample.EmployeeListener entity listener class , which has two life-cycle callback methods defined. The onListenerPrePersist method is for the PrePersist event, and the onListenerPostPersist method is for the

PostPersist event. Also, the checkEmployeeID method in the Employee class is configured to listen for the PrePersist event.

EntityManager fetch plan support

A FetchPlan is the strategy that the entity manager uses for retrieving associated objects if the application needs to access relationships.

Example

Assume for example that your application has two entities: Department and Employee. The relationship between the Department entity and the Employee entity is a bi-directional one-to-many relationship: One department has many employees, and one employee belongs to only one department. Since most of the time, when Department entity is fetched, its employees are likely to be fetched, the fetch type of this one-to-many relationship is set to be EAGER.

Here is a snippet of the Department class.

```
@Entity
public class Department {

    @Id
    private String deptId;

    @Basic
    String deptName;

    @OneToMany(fetch = FetchType.EAGER, mappedBy="department", cascade = {CascadeType.PERSIST})
    public Collection<Employee> employees;

}
```

In a distributed environment, when an application calls `em.find(Department.class, "dept1")` to find a Department entity with key "dept1", this find operation will get the Department entity and all its eager-fetched relations. In the case of the preceding snippet, these are all the employees of department "dept1".

Prior to WebSphere eXtreme Scale 6.1.0.5, the retrieval of one Department entity and N Employee entities incurred N+1 client-server trips because the client retrieved one entity for one client-server trip. You can improve performance if you retrieve these N+1 entities in one trip.

Fetch plan

A fetch plan can be used to customize how to fetch eager relationships by customizing the maximum depth of the relationships. The fetch depth overrides eager relations greater than the specified depth to lazy relations. By default, the fetch depth is the maximum fetch depth. This means that eager relationships of all levels that are eager-navigable from the root entity will be fetched. An EAGER relationship is eager-navigable from a root entity if and only if all the relations starting from the root entity to it are configured as eager-fetched.

In the previous example, the Employee entity is eager-navigable from the Department entity because the Department-Employee relationship is configured as eager-fetched.

If the Employee entity has another eager relationship to an Address entity for instance, then the Address entity is also eager-navigable from the Department entity. However, if the Department-Employee relationships were configured as

lazy-fetch, then the Address entity is not eager-navigable from the Department entity because the Department-Employee relationship breaks the eager fetch chain.

A FetchPlan object can be retrieved from the EntityManager instance. The application can use the setMaxFetchDepth method to change the maximum fetch depth.

A fetch plan is associated with an EntityManager instance. The fetch plan applies to any fetch operation, more specifically as follows.

- EntityManager find(Class class, Object key) and findForUpdate(Class class, Object key) operations
- Query operations
- QueryQueue operations

The FetchPlan object is mutable. Once changed, the changed value will be applied to the fetch operations executed afterward.

A fetch plan is important for a distributed deployment because it decides whether the eager-fetched relationship entities are retrieved with the root entity in one client-server trip or more than one.

Continuing with the previous example, consider further that the fetch plan has maximum depth set to infinity. In that case, when an application calls em.find(Department.class, "dept1") to find a Department, this find operation will get one Department entity and N employee entities in one client-server trip. However, for a fetch plan with maximum fetch depth set to zero, only the Department object will be retrieved from the server, while the Employee entities are retrieved from the server only when the employees collection of the Department object is accessed.

Different fetch plans

You have several different fetch plans based on your requirements, explained in the following sections.

Impact on a distributed grid

- *Infinite-depth fetch plan:* An infinite-depth fetch plan has its maximum fetch depth set to FetchPlan.DEPTH_INFINITE.

In a client-server environment, if an infinite-depth fetch plan is used, then all the relations that are eager-navigable from the root entity will be retrieved in one client-server trip.

Example: If the application is interested in all the Address entities of all employees of a particular Department, then it uses an infinite-depth fetch plan to retrieve all the associated Address entities. The following code only incurs one client-server trip.

```
em.getFetchPlan().setMaxFetchDepth(FetchPlan.DEPTH_INFINITE);

tran.begin();
Department dept = (Department) em.find(Department.class, "dept1");
// do something with Address object.
for (Employee e: dept.employees) {
    for (Address addr: e.addresses) {
        // do something with addresses.
    }
}
tran.commit();
```

- *Zero-depth fetch plan:* A zero-depth fetch plan has its maximum fetch depth set to 0.

In a client-server environment, if a zero fetch plan is used, then only the root entity will be retrieved in the first client-server trip. All the eager relationships are treated as if they were lazy.

Example: In this example, the application is only interested in the Department entity attribute. It does not need to access its employees, so the application sets the fetch plan depth to 0.

```
Session session = objectGrid.getSession();
EntityManager em = session.getEntityManager();
EntityTransaction tran = em.getTransaction();
em.getFetchPlan().setMaxFetchDepth(0);

tran.begin();
Department dept = (Department) em.find(Department.class, "dept1");
// do something with dept object.
tran.commit();
```

- *Fetch plan with depth k :*

A k -depth fetch plan has its maximum fetch depth set to k .

In a client-server eXtreme Scale environment, if a k -depth fetch plan is used, then all the relationships eager-navigable from the root entity within k steps will be retrieved in the first client-server trip.

The infinite-depth fetch plan ($k = \text{infinity}$) and zero-depth fetch plan ($k = 0$) are just two examples of the k -depth fetch plan.

To continue expanding on the previous example, assume there is another eager relationship from the entity Employee to the entity Address. If the fetch plan has maximum fetch depth set to 1, then the `em.find(Department.class, "dept1")` operation will retrieve the Department entity and all its Employee entities in one client-server trip. However, the Address entities will not be retrieved because they are not eager-navigable to the Department entity within 1 step, but 2 steps.

If you use a fetch plan with depth set to 2, then the `em.find(Department.class, "dept1")` operation will retrieve the Department entity, all its Employee entities, and all Address entities associated with the Employee entities in one client-server trip.

Tip: The default fetch plan has maximum fetch depth set to infinity, so the default behavior of a fetch operation can change. All the eager-navigable relationships from the root entity are retrieved. Instead of multiple trips, now the fetch operation only incurs one client-server trip with the default fetch plan. To keep the settings for the product from the prior version, set the fetch depth to 0.

- *Fetch plan used on query:*

If you execute an entity query you can also use a fetch plan to customize relationship retrieval.

For example, the query `SELECT d FROM Department d WHERE "d.deptName='Department'"` result has a relationship to the Department entity. Notice the fetch plan depth starts with the query result association: In this case, the Department entity, not the query result itself. That is, the Department entity is on fetch-depth level 0. Therefore a fetch plan with maximum fetch depth 1 will retrieve the Department entity and its Employee entities in one client-server trip.

Example: In this example, the fetch plan depth is set to 1, so the Department entity and its Employee entities are retrieved in one client-server trip, but the Address entities will not be retrieved in the same trip.

Important: If a relationship is ordered, using either `OrderBy` annotation or configuration, then it is considered an eager relationship even if it is configured as lazy-fetch.

Performance considerations in a distributed environment

By default, all relationships that are eager-navigable from the root entity will be retrieved in one client-server trip. This can improve performance if all the relationships are going to be used. However, in certain usage scenarios, not all relationships eager-navigable from the root entity are used, so they incur both run-time overhead and bandwidth overhead by retrieving those unused entities.

For such cases, the application can set the maximum fetch depth to a small number to decrease the depth of entities to be retrieved by making all the eager relations after that certain depth lazy. This setting can improve performance.

Proceeding still further with the previous Department-Employee-Address example, by default, all the Address entities associated with employees of the Department "dept1" will be retrieved when `em.find(Department.class, "dept1")` is called. If the application does not use Address entities, it can set the maximum fetch depth to 1, so the Address entities will not be retrieved with the Department entity.

EntityManager interface performance impact

An environment requiring every application to use the same data access objects for a given datastore would be highly impractical. In contrast, the EntityManager interface that is provided with WebSphere eXtreme Scale separates applications from the state held in its server grid datastore.

The cost of using the EntityManager interface is not high and depends on the type of work being performed. Always use the EntityManager interface and optimize the crucial business logic after the application is complete. You can rework any code that uses EntityManager interfaces to use maps and tuples. Generally, this code rework might be necessary for ten percent of the code.

If you use relationships between objects, then the performance impact is lower because an application that is using maps needs to manage those relationships similarly to the EntityManager interface.

Applications that use the EntityManager interface do not need to provide an ObjectTransformer because it is optimized automatically.

Reworking EntityManager code for maps

A sample entity follows:

```
@Entity
public class Person
{
    @Id
    String ssn;
    String firstName;
    @Index
    String middleName;
    String surname;
}
```

Some code to find the entity and update the entity follows:

```

Person p = null;
s.begin();
p = (Person)em.find(Person.class, "1234567890");
p.middleName = String.valueOf(inner);
s.commit();

```

The same code using Maps and Tuples follows:

```

Tuple key = null;
key = map.getEntityMetadata().getKeyMetadata().createTuple();
key.setAttribute(0, "1234567890");

// The Copy Mode is always NO_COPY for entity maps if not using COPY_TO_BYTES.
// Either we need to copy the tuple or we can ask the ObjectGrid to do it for us:
map.setCopyMode(CopyMode.COPY_ON_READ);
s.begin();
Tuple value = (Tuple)map.get(key);
value.setAttribute(1, String.valueOf(inner));
map.update(key, value);
value = null;
s.commit();

```

Both of these code snippets have the same result, and an application can use either or both snippets.

The second code snippet shows how to use maps directly and how to work with the tuples (the key and value pairs). The value tuple has three attributes: *firstName*, *middleName*, and *lastName*, indexed at 0, 1, and 2 respectively. The key tuple has a single attribute the ID number is indexed at zero. You can see how Tuples are created by using the `EntityMetadata#getKeyMetadata` or `EntityMetadata#getValueMetadata` methods. You must use these methods to create Tuples for an Entity. You cannot implement the Tuple interface and pass an instance of your Tuple implementation.

Instrumentation agent

You can improve the performance of field-access entities by enabling the WebSphere eXtreme Scale instrumentation agent when using Java Development Kit (JDK) Version 1.5 or later.

Enabling eXtreme Scale agent on JDK Version 1.5 or above

The ObjectGrid agent can be enabled with a Java command line option with the following syntax:

```
-javaagent:jarpath[=options]
```

The *jarpath* value is the path to an eXtreme Scale runtime Java archive (JAR) file that contains eXtreme Scale agent class and supporting classes such as the `objectgrid.jar`, `wsoobjectgrid.jar`, `ogclient.jar`, `wsogclient.jar`, and `ogagent.jar` files. Typically, in a stand-alone Java program or in a Java Platform, Enterprise Edition environment that is not running WebSphere Application Server, use the `objectgrid.jar` or `ogclient.jar` file. In a WebSphere Application Server or a multi-classloaders environment, you must use the `ogagent.jar` file in the Java command line agent option. Provide the `ogagent.config` file in the class path or use agent options to specify additional information.

eXtreme Scale agent options

config

Overrides the configuration file name.

include

Specifies or overrides transformation domain definition that is the first part of the configuration file.

exclude

Specifies or overrides the @Exclude definition.

fieldAccessEntity

Specifies or overrides the @FieldAccessEntity definition.

trace Specifies a trace level. Levels can be ALL, CONFIG, FINE, FINER, FINEST, SEVERE, WARNING, INFO, and OFF.

trace.file

Specifies the location of the trace file.

The semicolon (;) is used as a delimiter to separate each option. The comma (,) is used as a delimiter to separate each element within an option. The following example demonstrates the eXtreme Scale agent option for a Java program:

```
-javaagent:objectgridRoot/lib/objectgrid.jar=config=myConfigFile;
include=includedPackage;exclude=excludedPackage;
fieldAccessEntity=package1,package2
```

ogagent.config file

The ogagent.config file is the designated eXtreme Scale agent configuration file name. If the file name is in the class path, the eXtreme Scale agent finds and parses the file. You can override the designated file name through the config option of eXtreme Scale agent. The following example shows how to specify the configuration file:

```
-javaagent:objectgridRoot/lib/objectgrid.jar=config=myOverrideConfigFile
```

An eXtreme Scale agent configuration file has the following parts:

- **Transformation domain:** The transformation domain part is first in the configuration file. The transformation domain is a list of packages and classes that are included in the class transformation process. This transformation domain must include all classes that are field-access entity classes, and other classes that refer to these field-access entity classes. Field-access entity classes and those classes that refer to these field-access entity classes construct the transformation domain. If you plan to specify field-access entity classes in the @FieldAccessEntity part, then you do not need to include field-access entity classes here. The transformation domain must be complete. Otherwise, you might see a FieldAccessEntityNotInstrumentedException exception.
- **@Exclude:** The @Exclude token indicates that packages and classes listed after this token are excluded from the transformation domain.
- **@FieldAccessEntity:** The @FieldAccessEntity token indicates that packages and classes listed after this token are field-access Entity packages and classes. If no line exists after the @FieldAccessEntity token, then its equivalent is "No @FieldAccessEntity specified". The eXtreme Scale agent determines that there are no field-access Entity packages and classes defined. If there are lines after the @FieldAccessEntity token, then they represent the user-specified field-access Entity packages and classes. For example, "field-access entity domain". The field-access entity domain is a sub-domain of the transformation domain. Packages and classes that are listed in the field-access entity domain are a part of the transformation domain, even when they are not listed in the transformation domain. The @Exclude token, which lists packages and classes that are excluded from transformation, has no impact on the field-access Entity

domain. When `@FieldAccessEntity` token is specified, all field-access entities must be in this field-access Entity domain. Otherwise, a `FieldAccessEntityNotInstrumentedException` exception might occur.

Example agent configuration file (ogagent.config)

```
#####
# The # indicates comment line
#####
# This is an ObjectGrid agent config file (the designated file name is ogagent.config) that can be found and parsed by the ObjectGrid agent
# if it is in classpath.
# If the file name is "ogagent.config" and in classpath, Java program runs with -javaagent:objectgridRoot/ogagent.jar will have
# ObjectGrid agent enabled.
# If the file name is not "ogagent.config" but in classpath, you can specify the file name in config option of ObjectGrid agent
# -javaagent:objectgridRoot/lib/objectgrid.jar=config=myOverrideConfigFile
# See comments below for more info regarding instrumentation setting override.

# The first part of the configuration is the list of packages and classes that should be included in transformation domain.
# The includes (packages/classes, construct the instrumentation domain) should be in the beginning of the file.
com.testpackage
com.testClass

# Transformation domain: The above lines are packages/classes that construct the transformation domain.
# The system will process classes with name starting with above packages/classes for transformation.
#
# @Exclude token : Exclude from transformation domain.
# The @Exclude token indicates packages/classes after that line should be excluded from transformation domain.
# It is used when user want to exclude some packages/classes from above specified included packages
#
# @FieldAccessEntity token: Field-access Entity domain.
# The @FieldAccessEntity token indicates packages/classes after that line are field-access Entity packages/classes.
# If there is no line after the @FieldAccessEntity token, it is equivalent to "No @FieldAccessEntity specified".
# The runtime will consider the user does not specify any field-access Entity packages/classes.
# The "field-access Entity domain" is a sub-domain of transformation domain.
#
# Packages/classes listed in the "field-access Entity domain" will always be part of transformation domain,
# even they are not listed in transformation domain.
# The @Exclude, which lists packages/classes excluded from transformation, has no impact on the "field-access Entity domain".
# Note: When @FieldAccessEntity is specified, all field-access entities must be in this field-access Entity domain,
# otherwise, FieldAccessEntityNotInstrumentedException may occur.
#
# The default ObjectGrid agent config file name is ogagent.config
# The runtime will look for this file as a resource in classpath and process it.
# Users can override this designated ObjectGrid agent config file name via config option of agent.
#
# e.g.
# Javaagent:objectgridRoot/lib/objectgrid.jar=config=myOverrideConfigFile
#
# The instrumentation definition, including transformation domain, @Exclude, and @FieldAccessEntity can be overridden individually
# by corresponding designated agent options.
# Designated agent options include:
# include -> used to override instrumentation domain definition that is the first part of the config file
# exclude -> used to override @Exclude definition
# fieldAccessEntity -> used to override @FieldAccessEntity definition
#
# Each agent option should be separated by ":",
# Within the agent option, the package or class should be separated by "."
#
# The following is an example that does not override the config file name:
# -javaagent:objectgridRoot/lib/objectgrid.jar=include=includedPackage;exclude=excludedPackage;fieldAccessEntity=package1,package2
#####

@Exclude
com.excludedPackage
com.excludedClass

@FieldAccessEntity
```

Performance consideration

For better performance, specify the transformation domain and field-access entity domain.

Entity query queues

Query queues allow applications to create a queue qualified by a query in the server-side or local eXtreme Scale over an entity. Entities from the query result are stored in this queue. Currently, query queue is only supported in a map that is using the pessimistic lock strategy.

A query queue is shared by multiple transactions and clients. After the query queue becomes empty, the entity query associated with this queue is rerun and new results are added to the queue. A query queue is uniquely identified by the entity query string and parameters. There is only one instance for each unique query queue in one ObjectGrid instance. See the `EntityManager` API documentation for additional information.

Query queue example

The following example shows how query queue can be used.

```

/**
 * Get a unassigned question type task
 */
private void getUnassignedQuestionTask() throws Exception {
    EntityManager em = og.getSession().getEntityManager();
    EntityTransaction tran = em.getTransaction();

    QueryQueue queue = em.createQueryQueue("SELECT t FROM Task t
    WHERE t.type=?1 AND t.status=?2", Task.class);
    queue.setParameter(1, new Integer(Task.TYPE_QUESTION));
    queue.setParameter(2, new Integer(Task.STATUS_UNASSIGNED));

    tran.begin();
    Task nextTask = (Task) queue.getNextEntity(10000);
    System.out.println("next task is " + nextTask);
    if (nextTask != null) {
        assignTask(em, nextTask);
    }
    tran.commit();
}

```

The previous example first creates a QueryQueue with a entity query string, "SELECT t FROM Task t WHERE t.type=?1 AND t.status=?2". Then it sets the parameters for the QueryQueue object. This query queue represents all "unassigned" tasks of the type "question". The QueryQueue object is very similar to an entity Query object.

After the QueryQueue is created, an entity transaction is started and the getNextEntity method is invoked, which retrieves the next available entity with a timeout value set to 10 seconds. After the entity is retrieved, it is processed in the assignTask method. The assignTask modifies the Task entity instance and changes the status to "assigned" which effectively removes it from the queue since it no longer matches the QueryQueue's filter. Once assigned, the transaction is committed.

From this simple example, you can see a query queue is similar to an entity query. However, they differ in the following ways:

1. Entities in the query queue can be retrieved in an iterative manner. The user application decides the number of entities to be retrieved. For example, if QueryQueue.getNextEntity(timeout) is used, only one entity is retrieved, and if QueryQueue.getNextEntities(5, timeout) is used, 5 entities are retrieved. In a distributed environment, the number of entities directly decides the number of bytes to be transferred from the server to client.
2. When an entity is retrieved from the query queue, a U lock is placed on the entity so no other transactions can access it.

Retrieve entities in a loop

You can retrieve entities in a loop. An example that illustrates how to get all the unassigned, question type tasks completed follows.

```

/**
 * Get all unassigned question type tasks
 */
private void getAllUnassignedQuestionTask() throws Exception {
    EntityManager em = og.getSession().getEntityManager();
    EntityTransaction tran = em.getTransaction();

    QueryQueue queue = em.createQueryQueue("SELECT t FROM Task t WHERE
    t.type=?1 AND t.status=?2", Task.class);
    queue.setParameter(1, new Integer(Task.TYPE_QUESTION));

```

```

queue.setParameter(2, new Integer(Task.STATUS_UNASSIGNED));

Task nextTask = null;

do {
    tran.begin();
    nextTask = (Task) queue.getNextEntity(10000);
    if (nextTask != null) {
        System.out.println("next task is " + nextTask);
    }
    tran.commit();
} while (nextTask != null);
}

```

If there are 10 unassigned question-type tasks in the entity map, you might expect that you will have 10 entities printed to the console. However, if this example is run, you will see the program never exits, which might be contrary to what you assumed.

When a query queue is created and the `getNextEntity` is called, the entity query associated with the queue is executed and the 10 results are populated into the queue. When `getNextEntity` is called, an entity is taken off the queue. After 10 `getNextEntity` calls are executed, the queue is empty. The entity query will automatically re-run. Since these 10 entities still exist and match the query queue's filter criteria, they are populated into the queue again.

If the following line is added after the `println()` statement, you will see only 10 entities printed.

```
em.remove(nextTask);
```

For information on using `SessionHandle` with `QueryQueue` in a per-container placement deployment, read about `SessionHandle` integration.

Query queues deployed to all partitions

In a distributed eXtreme Scale, a query queue can be created for one partition or all partitions. If a query queue is created for all partitions, there will be one query queue instance in each partition.

When a client tries to get the next entity using the `QueryQueue.getNextEntity` or `QueryQueue.getNextEntities` method, the client sends a request to one of the partitions. A client sends peek and pin requests to the server:

- With a peek request, the client sends a request to one partition and the server returns immediately. If there is an entity in the queue, the server sends a response with the entity; if there is not, the server sends a response with no entity. In either case, the server will return immediately.
- With a pin request, the client sends a request to one partition and the server waits until an entity is available. If there is an entity in the queue, the server sends a response with the entity immediately; if there is not, the server waits on the queue until either an entity is available or the request times out.

An example of how an entity is retrieved for a query queue which is deployed to all partitions (n) follows:

1. When a `QueryQueue.getNextEntity` or `QueryQueue.getNextEntities` method is called, the client picks a random partition number from 0 to n-1.
2. The client sends peek request to the random partition.

- If an entity is available, the `QueryQueue.getNextEntity` or `QueryQueue.getNextEntities` method exits by returning the entity.
- If an entity is not available and is not the last unvisited partition, the client sends a peek request to the next partition.
- If an entity is not available and it is the last unvisited partition, the client instead sends a pin request.
- If the pin request to the last partition times-out and there is still no data available, the client will make a last effort by sending peek request to all partitions serially one more round. Therefore, if any entity is available in the previous partitions, the client will be able to get it.

Subset entity and no-entity support

The method to create a `QueryQueue` object in the entity manager follows:

```
public QueryQueue createQueryQueue(String qlString, Class entityClass);
```

The result in the query queue should be projected to the object defined by the second parameter to the method, `Class entityClass`.

If this parameter is specified, the class must have the same entity name as specified in the query string. This is useful if you want to project an entity into a subset entity. If a null value is used as the entity class, then the result will not be projected. The value stored in the map will be in an entity tuple format.

Client-side key collision

In distributed eXtreme Scale environment, query queue is only supported for eXtreme Scale maps with pessimistic locking mode. Therefore, there is no near cache on the client side. However, a client could have data (key and value) in the transactional map. This potentially could lead to a key collision when an entity retrieved from the server share the same key as an entry already in the transactional map.

When a key collision happens, the eXtreme Scale client run time uses the following rule to either throw an exception or silently override the data.

1. If the collided key is the key of the entity specified in the entity query associated with the query queue, then an exception is thrown. In this case, the transaction is rolled back, and the U lock on this entity key will be released on the server side.
2. Otherwise, if the collided key is the key of the entity association, the data in the transactional map will be overridden without warning.

The key collision only happens when there is a data in the transactional map. In other words, it only happens when a `getNextEntity` or `getNextEntities` call is called in a transaction which has already been dirtied (a new data has been inserted or a data has been updated). If an application does not want a key collision happen, it should always call `getNextEntity` or `getNextEntities` in a transaction which has not been dirtied.

Client failures

After a client sends a `getNextEntity` or `getNextEntities` request to the server, the client could fail as follows:

1. The client sends a request to the server and then goes down.
2. The client gets one or more entities from the server and then goes down.

In the first case, the server discovers that the client is going down when it tries to send back the response to the client. In the second case, when the client gets one or more entities from the server, an X lock is placed on these entities. If the client goes down, the transaction will eventually time out, and the X lock will be released.

Query with ORDER BY clause

Generally, query queues do not honor the ORDER BY clause. If you call `getNextEntity` or `getNextEntities` from the query queue, there is no guarantee the entities are returned according to the order. The reason is that the entities cannot be ordered across partitions. In the case that the query queue is deployed to all partitions, when a `getNextEntity` or `getNextEntities` call is executed, a random partition is picked to process the request. Therefore, the order is not guaranteed.

ORDER BY is honored if a query queue is deployed to a single partition.

For more information see “EntityManager Query API” on page 102.

One call per transaction

Each `QueryQueue.getNextEntity` or `QueryQueue.getNextEntities` call retrieves matched entities from one random partition. Applications should call exactly one `QueryQueue.getNextEntity` or `QueryQueue.getNextEntities` on one transaction. Otherwise eXtreme Scale could end up touching entities from multiple partitions, causing an exception to be thrown at the commit time.

EntityTransaction interface

You can use the EntityTransaction interface to demarcate transactions.

Purpose

To demarcate a transaction, you can use the EntityTransaction interface, which is associated with an entity manager instance. Use the `EntityManager.getTransaction` method to retrieve the EntityTransaction instance for the entity manager. Each `EntityManager` and `EntityTransaction` instance are associated with the `Session`. You can demarcate transactions with either the `EntityTransaction` or `Session`. Methods on the EntityTransaction interface do not have any checked exceptions. Only runtime exceptions of type `PersistenceException` or its subclasses result.

For more information about the EntityTransaction interface, see the API documentationEntityTransaction interface in the API documentation.

Entity manager tutorial: Overview

The tutorial for the entity manager shows you how to use WebSphere eXtreme Scale to store order information on a Web site. You can create a simple Java Platform, Standard Edition 5 application that uses an in-memory, local eXtreme Scale. The entities use Java SE 5 annotations and generics.

Before you begin

Ensure that you have met the following requirements before you begin the tutorial:

- You must have Java SE 5.
- You must have the `objectgrid.jar` file in your classpath.

Retrieving entities and objects (Query API)

WebSphere eXtreme Scale provides a flexible query engine for retrieving entities using the EntityManager API and Java objects using the ObjectQuery API.

WebSphere eXtreme Scale query capabilities

With the eXtreme Scale query engine, you can perform SELECT type queries over an entity or object-based schema using the eXtreme Scale query language.

This query language provides the following capabilities:

- Single and multi-valued results
- Aggregate functions
- Sorting and grouping
- Joins
- Conditional expressions with subqueries
- Named and positional parameters
- eXtreme Scale index use
- Path expression syntax for object navigation
- Pagination

Query interface

Use the query interface to control entity query execution.

Use the `EntityManager.createQuery(String)` method to create a Query. You can use each query instance multiple times with the EntityManager instance in which it was retrieved.

Each query result produces an entity, where the entity key is the row ID (of type long) and the entity value contains the field results of the SELECT clause. You can use each query result in subsequent queries.

The following methods are available on the `com.ibm.websphere.objectgrid.em.Query` interface.

public ObjectMap getResultMap()

The `getResultMap` method runs a SELECT query and returns the results in an ObjectMap object with the results in query-specified order. The resulting ObjectMap is valid only for the current transaction.

The map key is the result number, of type long, starting at 1. The map value is of type `com.ibm.websphere.projector.Tuple` where each attribute and association is named based on its ordinal position within the select clause of the query. Use the method to retrieve the EntityMetadata for the Tuple object that is stored within the map.

The `getResultMap` method is the fastest method for retrieving query result data where multiple results can exist. You can retrieve the name of the resulting entity using the `ObjectMap.getEntityMetadata()` and `EntityMetadata.getName()` methods.

Example: The following query returns two rows.

```
String q1 = SELECT e.name, e.id, d from Employee e join e.dept d WHERE d.number=5
Query q = em.createQuery(q1);
ObjectMap resultMap = q.getResultMap();
long rowID = 1; // starts with index 1
Tuple tResult = (Tuple) resultMap.get(new Long(rowID));
while(tResult != null) {
    // The first attribute is name and has an attribute name of 1
    // But has an ordinal position of 0.
    String name = (String)tResult.getAttribute(0);
    Integer id = (String)tResult.getAttribute(1);

    // Dept is an association with a name of 3, but
    // an ordinal position of 0 since it's the first association.
    // The association is always a OneToOne relationship,
    // so there is only one key.
    Tuple deptKey = tResult.getAssociation(0,0);
    ...
    ++rowID;
    tResult = (Tuple) resultMap.get(new Long(rowID));
}
}
```

public Iterator getResultIterator

The `getResultIterator` method runs a `SELECT` query and returns the query results using an `Iterator` where each result is either an `Object` for a single-valued query, or an `Object` array for a multiple-valued query. The values in the `Object[]` result are stored in query order. The result `Iterator` is valid for the current transaction only.

This method is preferred for retrieving query results within the `EntityManager` context. You can use the optional `setResultEntityName(String)` method to name the resulting entity so that it can be used in further queries.

Example: The following query returns two rows.

```
String q1 = SELECT e.name, e.id, e.dept from Employee e WHERE e.dept.number=5
Query q = em.createQuery(q1);
Iterator results = q.getResultIterator();
while(results.hasNext()) {
    Object[] curEmp = (Object[]) results.next();
    String name = (String) curEmp[0];
    Integer id = (Integer) curEmp[1];
    Dept d = (Dept) curEmp[2];
    ...
}
}
```

public Iterator getResultIterator(Class resultType)

The `getResultIterator(Class resultType)` method runs a `SELECT` query and returns the query results using an entity `Iterator`. The entity type is determined by the `resultType` parameter. The result `Iterator` is valid only for the current transaction.

Use this method when you want to use the `EntityManager` APIs to access the resulting entities.

Example: The following query returns all of the employees and the department to which they belong for one division, ordering by salary. To print out the five employees with the highest salaries and then select work with employees from only one department in the same working set, use the following code.

```
String string_q1 = "SELECT e.name, e.id, e.dept from Employee e WHERE
    e.dept.division='Manufacturing' ORDER BY e.salary DESC";
Query query1 = em.createQuery(string_q1);
query1.setResultEntityName("AllEmployees");
Iterator results1 = query1.getResultIterator(EmployeeResult.class);
```

```

int curEmployee = 0;
System.out.println("Highest paid employees");
while (results1.hasNext() && curEmployee++ < 5) {
    EmployeeResult curEmp = (EmployeeResult) results1.next();
    System.out.println(curEmp);
    // Remove the employee from the resultset.
    em.remove(curEmp);
}

// Flush the changes to the result map.
em.flush();

// Run a query against the local working set without the employees we
// removed
String string_q2 = "SELECT e.name, e.id, e.dept from AllEmployees e
    WHERE e.dept.name='Hardware'";
Query query2 = em.createQuery(string_q2);
Iterator results2 = query2.getResultIterator(EmployeeResult.class);
System.out.println("Subset list of Employees");
while (results2.hasNext()) {
    EmployeeResult curEmp = (EmployeeResult) results2.next();
    System.out.println(curEmp);
}

```

public Object getSingleResult

The `getSingleResult` method runs a SELECT query that returns a single result.

If the SELECT clause has more than one field defined, then the result is an object array, where each element in the array is based on its ordinal position within the SELECT clause of the query.

```

String q1 = "SELECT e from Employee e WHERE e.id=100"
Employee e = em.createQuery(q1).getSingleResult();

String q1 = "SELECT e.name, e.dept from Employee e WHERE e.id=100"
Object[] empData = em.createQuery(q1).getSingleResult();
String empName= (String) empData[0];
Department empDept = (Department) empData[1];

```

public Query setResultEntityName(String entityName)

The `setResultEntityName(String entityName)` method specifies the name of the query result entity.

Each time the `getResultIterator` or `getResultMap` methods are invoked, an entity with an `ObjectMap` is dynamically created to hold the results of the query. If the entity is not specified, or null, the entity and `ObjectMap` name are automatically generated.

Because all query results are available for the duration of a transaction, a query name cannot be reused in a single transaction.

public Query setPartition(int partitionId)

Set the partition to where the query routes.

This method is required if the maps in the query are partitioned and if the entity manager does not have affinity to a single schema root entity partition.

Use the `PartitionManager` Interface to determine the number of partitions for the backing map of a given entity.

The following table provides descriptions of the other methods that are available through the query interface.

Table 2. Other methods.

| Method | Result |
|---|--|
| public Query setMaxResults(int maxResult) | Set the maximum number of results to retrieve. |
| public Query setFirstResult(int startPosition) | Set the position of the first result to retrieve. |
| public Query setParameter(String name, Object value) | Bind an argument to a named parameter. |
| public Query setParameter(int position, Object value) | Bind an argument to a positional parameter. |
| public Query setFlushMode(FlushModeType flushMode) | Set the flush mode type to be used when the query runs, overriding the flush mode type set on the EntityManager. |

eXtreme Scale query elements

With the eXtreme Scale query engine, you can use a single query language for searching the eXtreme Scale cache. This query language can query Java objects that are stored in ObjectMap objects and Entity objects. Use the following syntax for creating a query string.

An eXtreme Scale query is a string that contains the following elements:

- A SELECT clause that specifies the objects or values to return.
- A FROM clause that names the object collections.
- An optional WHERE clause that contains search predicates over the collections.
- An optional GROUP BY and HAVING clause (see eXtreme Scale query aggregation functions).
- An optional ORDER BY clause that specifies the ordering of the result collection.

Collections of Java objects are identified in queries through the use of their name in the query FROM clause.

The elements of query language are discussed in more detail in the following related topics:

- “ObjectGrid query Backus-Naur Form” on page 114 syntax
- “Reference for eXtreme Scale queries” on page 106

The following topics describe the means to use the Query API:

- “EntityManager Query API” on page 102
- “Using the ObjectQuery API” on page 97

Querying data in multiple time zones

In a distributed scenario, queries actually run on servers. When querying data with predicates of type calendar, java.util.Date and timestamp, the specified date time value in a query is based on the local time zone of the server.

In a single time-zone system where all clients and servers run on same time zone, you do not need to consider issues related to predicate types with calendar, java.util.Date and timestamp. However, when clients and servers are in different

time zones, the specified date time value in queries is based on the server time zone and may return unwanted data back to client. Without knowing the server time zone, the specified date time value is meaningless. So the specified date time value should consider the time zone offset difference between the target time zone and the server time zone.

Time zone offset

For example, assume that a client is in [GMT-0] time zone and the server is in [GMT-6] time zone. The server time zone is 6 hours behind the client. The client would like to run the following query:

```
SELECT e FROM Employee e WHERE e.birthDate='1999-12-31 06:00:00'
```

Assuming the entity Employee has a birthDate attribute that is of type `java.util.Date`, the client is in [GMT-0] time zone and wants to retrieve Employees with birthDate value as '1999-12-31 06:00:00 [GMT-0]' based on its time zone.

The query will run on the server and the birthDate value used by the query engine will be '1999-12-31 06:00:00 [GMT-6]' that equals to '1999-12-31 12:00:00 [GMT-0]'. Employees with birthDate value equal to '1999-12-31 12:00:00 [GMT-0]' will be returned to the client. Thus, the client will not get wanted Employees with birthDate value '1999-12-31 06:00:00 [GMT-0]'.

The problem described occurs because of the time zone difference between client and server. To solve this problem, one approach is to calculate the time zone offset between client and server and apply the time zone offset on the target date time value in the query. In the previous query example, the time zone offset is -6 hours, and the adjusted birthDate predicate should be "birthDate='1999-12-31 00:00:00'" if the client intends to retrieve Employees with birthDate value '12-31 06:00:00 [GMT-0]'. With the adjusted birthDate value, the server will use '1999-12-31 00:00:00 [GMT-6]' that equals to target value '12-31 06:00:00 [GMT-0]', and the required Employees will be returned to the client.

Distributed deployment in multiple time zones

If the distributed eXtreme Scale grid is deployed into multiple ObjectGrid servers in various time zones, the adjusting time zone offset approach will not work because the client will not know which server will run the query and thus cannot determine the time zone offset to use. The only solution is to use suffix 'Z' (not case sensitive) on JDBC date and time escape format to indicate using GMT time zone based date time value. The suffix 'Z' (not case sensitive) indicates to use GMT time zone based date time value. Without the suffix 'Z', the local time zone based date time value will be used in the process that runs the query.

The following query is equivalent to the previous example, but uses the suffix 'Z' instead:

```
SELECT e FROM Employee e WHERE e.birthDate='1999-12-31 06:00:00Z'
```

The query should find Employees with birthDate value '1999-12-31 06:00:00'. The suffix 'Z' indicates the specified birthDate value is GMT time zone based, so the GMT time zone based birthDate value '1999-12-31 06:00:00 [GMT-0]' will be used by the query engine for matching criteria value. Employees with birthDate attribute value equal to this GMT based birthDate value '1999-12-31 06:00:00 [GMT-0]' will be included in query result. Using the suffix 'Z' on JDBC date time escape format in any query is crucial to make applications time zone safe. Without

this approach, the date time value is server time zone based and is meaningless from the client perspective when clients and servers are in different time zones.

For more information, see the topic on inserting data for different time zones in the *Product Overview*.

Inserting data for different time zones

When inserting data with `calendar`, `java.util.Date`, and `timestamp` attributes into an `ObjectGrid`, you must ensure these date time attributes are created based on same time zone, especially when deployed into multiple servers in various time zones. Using the same time zone based date time objects can ensure the application is time-zone safe and data can be queried by `calendar`, `java.util.Date` and `timestamp` predicates.

Without explicitly specifying a time zone when creating date time objects, Java will use the local time zone and may cause inconsistent date time values in clients and servers.

Consider an example in a distributed deployment in which `client1` is in time zone `[GMT-0]` and `client2` is in `[GMT-6]` and both want to create a `java.util.Date` object with value `'1999-12-31 06:00:00'`. Then `client1` will create `java.util.Date` object with value `'1999-12-31 06:00:00 [GMT-0]'` and `client2` will create `java.util.Date` object with value `'1999-12-31 06:00:00 [GMT-6]'`. Both `java.util.Date` objects are not equal because the time zone is different. A similar problem occurs when preloading data into partitions residing in servers in different time zones if local time zone is used to create date time objects.

To avoid the described problem, the application can choose a time zone such as `[GMT-0]` as the base time zone for creating `calendar`, `java.util.Date`, and `timestamp` objects.

For more information, see the topic on querying data in multiple time zones in the *Programming Guide*.

Using the ObjectQuery API

The `ObjectQuery` API provides methods for querying data in the `ObjectGrid` that is stored using the `ObjectMap` API. When a schema is defined in the `ObjectGrid` instance, the `ObjectQuery` API can be used to create and run queries over the heterogeneous objects stored in the object maps.

Query and object maps

You can use an enhanced query capability for objects that are stored using the `ObjectMap` API. These queries allow retrieval of objects using non-key attributes and performs simple aggregations such as `sum`, `avg`, `min`, and `max` against all the data that matches a query. Applications can construct a query using the `Session.createObjectQuery` method. This method returns an `ObjectQuery` object which can then be interrogated to obtain the query results. The query object also allows the query to be customized before running the query. The query is run automatically when any method returning the result is called.

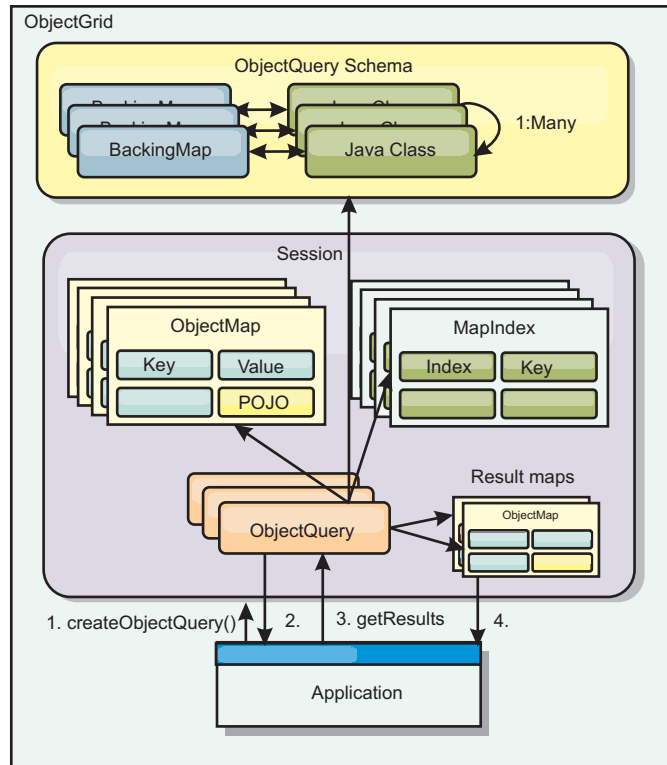


Figure 1. The interaction of the query with the ObjectGrid object maps and how a schema is defined for classes and associated with an ObjectGrid map

Defining an ObjectMap schema

Object maps are used to store objects in various forms and are largely unaware of the format. A schema must be defined in the ObjectGrid that defines the format of the data. A schema is composed of the following pieces:

- The type of object stored in the ObjectMap
- Relationships between ObjectMaps
- The method for which each query should access the data attributes in the objects (fields or property methods)
- The primary key attribute name in the object.

See *Configuring an ObjectQuery schema* for details.

For an example on creating a schema programmatically or using the ObjectGrid descriptor XML file, see the tutorial on the ObjectQuery in the *Product Overview*.

Querying objects with the ObjectQuery API

The ObjectQuery interface allows the querying of non-entity objects, which are heterogeneous objects that are stored directly in the ObjectGrid ObjectMaps. The ObjectQuery API provides an easy way to find ObjectMap objects without using the keyword and index mechanisms directly.

There are two methods for retrieving results from an ObjectQuery: `getResultIterator` and `getResultMap`.

Retrieving query results using getResultIterator

Query results are basically a list of attributes. Suppose the query was `select a,b,c from X where y=z`. This query returns a list of rows containing a, b and c. This list is actually stored in a transaction scoped Map, which means that you must associate an artificial key with each row and use an integer that increases with each row. This map is obtained using the `ObjectQuery.getResultMap()` method. You can access the elements of each row using code similar to the following:

```
ObjectQuery q = session.createQuery(
    "select c.id, c.firstName, c.surname from Customer c where c.surname=?1");

q.setParameter(1, "Claus");

Iterator iter = q.getResultIterator();
while(iter.hasNext())
{
    Object[] row = (Object[])iter.next();
    System.out.println("Found a Claus with id "
        + row[objectgrid: 0 ] + ", firstName: "
        + row[objectgrid: 1 ] + ", surname: "
        + row[objectgrid: 2 ]);
}
```

Retrieving query results using getResultMap

Query results can also be retrieved using the result map directly. The following example shows a query retrieving specific parts of the matching Customers and demonstrates how to access the resulting rows. Notice that if you use the `ObjectQuery` object to access the data, then the generated long row identifier is hidden. The long row is only visible when using the `ObjectMap` to access the result.

When the transaction is completed this map disappears. The map is also only visible to the session used, that is, normally to just the thread that created it. The map uses a key of type `Long` which represents the row ID. The values stored in the map either are of type `Object` or `Object[]`, where each element matches the type of the element in the select clause of query.

```
ObjectQuery q = em.createQuery(
    "select c.id, c.firstName, c.surname from Customer c where c.surname=?1");
q.setParameter(1, "Claus");
ObjectMap qmap = q.getResultMap();
for(long rowId = 0; true; ++rowId)
{
    Object[] row = (Object[]) qmap.get(new Long(rowId));
    if(row == null) break;
    System.out.println(" I Found a Claus with id " + row[0]
        + ", firstName: " + row[1]
        + ", surname: " + row[2]);
}
```

For examples on using the `ObjectQuery`, see the tutorial on the `ObjectQuery` API in the *Product Overview*.

Configuring an ObjectQuery schema

`ObjectQuery` relies on schema or shape information to perform semantic checking and to evaluate path expressions. This section describes how to define the schema in XML or programmatically.

Defining the schema

The ObjectMap schema is defined in the ObjectGrid deployment descriptor XML or programmatically using the normal eXtreme Scale configuration techniques. For an example on how to create a schema, see “Configuring an ObjectQuery schema” on page 99

Schema information describes plain old Java objects (POJOs): which attributes they consist of and what types of attributes there might be, whether the attributes are primary key fields, single-valued or multi-valued relationships, or bidirectional relationships. Schema information directs ObjectQuery to use field access or property access.

Queryable attributes

When the schema is defined in the ObjectGrid, the objects in the schema are introspected using reflection to determine which attributes are available for querying. You can query the following attribute types:

- Java primitive types including wrappers
- java.lang.String
- java.math.BigInteger
- java.math.BigDecimal
- java.util.Date
- java.sql.Date
- java.sql.Time
- java.sql.Timestamp
- java.util.Calendar
- byte[]
- java.lang.Byte[]
- char[]
- java.lang.Character[]
- J2SE enum

Embedded serializable types other than those stated previously can also be included in a query result, but cannot be included in the WHERE or FROM clause of the query. Serializable attributes are not navigable.

Attribute types can be excluded from the schema if the type is not serializable, the field or property is static, or the field is transient. Since all map objects must be serializable, the ObjectGrid only includes attributes that can be persisted from the object. Other objects are ignored.

Field attributes

When the schema is configured to access the object using fields, all serializable, non-transient fields are automatically incorporated into the schema. To select a field attribute in a query, use the field identifier name as it exists in the class definition.

All public, private, protected and package protected fields are included in the schema.

Property attributes

When the schema is configured to access the object using properties, all serializable methods that follow the JavaBeans property naming conventions will automatically be incorporated into the schema. To select a property attribute for the query, use the JavaBeans style property name conventions.

All public, private, protected and package protected properties are included in the schema.

In the following class, the following attributes are added to the schema: name, birthday, valid.

```
public class Person {
    public String getName(){}
    private java.util.Date getBirthday(){}
    boolean isValid(){}
    public NonSerializableObject getData(){}
}
```

When using a CopyMode of COPY_ON_WRITE, the query schema must always use property-based access. COPY_ON_WRITE creates proxy objects whenever objects are retrieved from the map and can only access those objects using property methods. Failure to do so will result in each query result being set to null.

Relationships

Each relationship must be explicitly defined in the schema configuration. The cardinality of the relationship is automatically determined by the type of the attribute. If the attribute implements the java.util.Collection interface, then the relationship is either a one-to-many or many-to-many relationship.

Unlike entity queries, attributes that refer to other cached objects must not store direct references to the object. References to other objects are serialized as part of the containing object's data. Instead, store the key to the related object.

For example, if there is a many-to-one relationship between a Customer and Order:
Incorrect. Storing an object reference.

```
public class Customer {
    String customerId;
    Collection<Order> orders;
}
```

```
public class Order {
    String orderId;
    Customer customer;
}
```

Correct. The key to the related object.

```
public class Customer {
    String customerId;
    Collection<String> orders;
}
```

```
public class Order {
    String orderId;
    String customer;
}
```

When a query is run that joins the two map objects together, the key will automatically be inflated. For example, the following query would return Customer objects:

```
SELECT c FROM Order o JOIN Customer c WHERE orderId=5
```

Using indexes

ObjectGrid uses index plugins to add indexes to maps. The query engine automatically incorporates any indexes that are defined on a schema map element of the type: `com.ibm.websphere.objectgrid.plugins.index.HashIndex` and the `rangeIndex` property is set to true. If the index type is not `HashIndex` and the `rangeIndex` property is not set to true, then the index is ignored by the query. See Object Query tutorial the `ObjectQuery` tutorial in the *Product Overview* for an example on how to add an index to the schema.

EntityManager Query API

The EntityManager API provides methods for querying data in the ObjectGrid that is stored using the EntityManager API. The EntityManager Query API is used to create and run queries over one or more entities defined in eXtreme Scale.

Query and ObjectMaps for entities

WebSphere Extended Deployment v6.1 introduced an enhanced query capability for entities stored in eXtreme Scale. These queries allow objects to be retrieved using non-key attributes and to perform simple aggregations such as the sum, average, minimum, and maximum against all the data that matches a query. Applications construct a query using the `EntityManager.createQuery` API. This returns a Query object and can then be interrogated to obtain the query results. The query object also allows the query to be customized before running the query. The query is run automatically when any method returning the result is called.

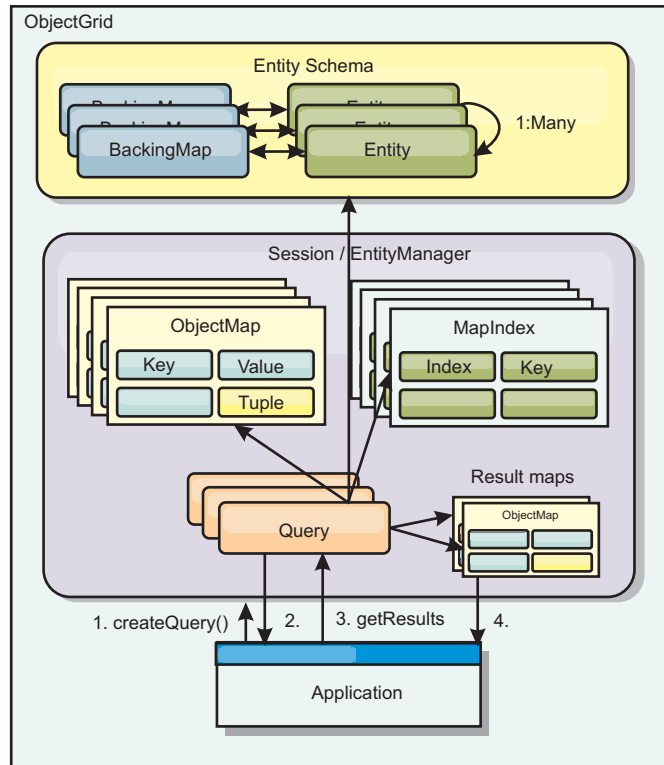


Figure 2. The interaction of the query with the ObjectGrid object maps and how the entity schema is defined and associated with an ObjectGrid map.

Retrieving query results using the getResultIterator method

Query results are a list of attributes. If the query was `select a,b,c from X where y=z`, then a list of rows containing `a`, `b` and `c` is returned. This list is stored in a transaction scoped Map, which means that you must associated an artificial key with each row and use an integer that increases with each row. This map is obtained using the `Query.getResultMap` method. The map has `EntityMetaData`, which describes each row in the Map associated with it. You can access the elements of each row using code similar to the following:

```
Query q = em.createQuery("select c.id, c.firstName, c.surname from Customer c where c.surname=?1");
q.setParameter(1, "Claus");

Iterator iter = q.getResultIterator();
while(iter.hasNext())
{
    Object[] row = (Object[])iter.next();
    System.out.println("Found a Claus with id " + row[objectgrid: 0 ]
        + ", firstName: " + row[objectgrid: 1 ]
        + ", surname: " + row[objectgrid: 2 ]);
}
```

Retrieving query results using getResultMap

The following code shows the retrieval of specific parts of the matching Customers and shows how to access the resulting rows. If you use the Query object to access the data, then the generated long row identifier is hidden. The long is only visible when using the ObjectMap to access the result. When the transaction is completed, then this Map disappears. The Map is only visible to the Session used, that is, normally to just the thread that created it. The Map uses a Tuple for the key with a single attribute, a long with the row ID. The value is another tuple with an attribute for each column in the result set.

The following sample code demonstrates this:

```
Query q = em.createQuery("select c.id, c.firstName, c.surname from
Customer c where c.surname=?1");
q.setParameter(1, "Claus");
ObjectMap qmap = q.getResultMap();
Tuple keyTuple = qmap.getEntityMetadata().getKeyMetadata().createTuple();
for(long i = 0; true; ++i)
{
    keyTuple.setAttribute(0, new Long(i));
    Tuple row = (Tuple)qmap.get(keyTuple);
    if(row == null) break;
    System.out.println(" I Found a Claus with id " + row.getAttribute(0)
        + ", firstName: " + row.getAttribute(1)
        + ", surname: " + row.getAttribute(2));
}
```

Retrieving query results using an entity result iterator

The following code shows the query and the loop that retrieves each result row using the normal Map APIs. The key for the Map is a Tuple. So, construct one of the correct types using the createTuple method result in keyTuple. Try to retrieve all rows with rowIds from 0 onwards. When you get returns null (indicating key not found), then the loop finishes. Set the first attribute of keyTuple to be the long that you want to find. The value returned by get is also a Tuple with an attribute for each column in the query result. Then, pull each attribute from the value Tuple using getAttribute.

Following is the next code fragment:

```
Query q2 = em.createQuery("select c.id, c.firstName, c.surname from Customer c where c.surname=?1");
q2.setResultEntityName("CustomerQueryResult");
q2.setParameter(1, "Claus");

Iterator iter2 = q2.getResultIterator(CustomerQueryResult.class);
while(iter2.hasNext())
{
    CustomerQueryResult row = (CustomerQueryResult)iter2.next();
    // firstName is the id not the firstName.
    System.out.println("Found a Claus with id " + row.id
        + ", firstName: " + row.firstName
        + ", surname: " + row.surname);
}

em.getTransaction().commit();
```

Specified is a ResultEntityName value on the query. This value tells the query engine that you want to project each row to a specific object, CustomerQueryResult in this case. The class follows:

```
@Entity
public class CustomerQueryResult {
    @Id long rowId;
    String id;
    String firstName;
    String surname;
};
```

In the first snippet, notice that the each query row is returned as a CustomerQueryResult object rather than an Object[]. The result columns of the query are projected to the CustomerQueryResult object. Projecting the result is slightly slower at run time but more readable. Query result Entities should not be registered with eXtreme Scale at startup. If the entities are registered, then a global Map with the same name is created, and the query fails with an error indicating duplicate Map name.

Simple queries with EntityManager

WebSphere eXtreme Scale comes with EntityManager query API.

The EntityManager query API is very similar to SQL other query engines that query over objects. A query is defined, then the result is retrieved from the query using various getResult methods.

The following examples refer to the entities used in the EntityManager tutorial in the Product Overview.

Running a simple query

In this example, customers with the surname of Claus are queried:

```
em.getTransaction().begin();

Query q = em.createQuery("select c from Customer c where c.surname='Claus'");

Iterator iter = q.getResultIterator();
while(iter.hasNext())
{
    Customer c = (Customer)iter.next();
    System.out.println("Found a claus with id " + c.id);
}

em.getTransaction().commit();
```

Using parameters

Since you want to find all customers with a surname of Claus, a parameter to specify the surname is used since you might want to use this query more than once.

Positional Parameter Example

```
Query q = em.createQuery("select c from Customer c where c.surname=?1");
q.setParameter(1, "Claus");
```

Using parameters is very important when the query is used more than once. The EntityManager needs to parse the query string and build a plan for the query, which is expensive. By using a parameter, the EntityManager caches the plan for the query, thereby reducing the time it takes to run a query.

Both positional and named parameters are used:

Named Parameter Example

```
Query q = em.createQuery("select c from Customer c where c.surname=:name");
q.setParameter("name", "Claus");
```

Using an index to improve performance

If there are millions of customers, then the previous query needs to scan over all rows in the Customer Map. This is not very efficient. But eXtreme Scale provides a mechanism for defining indexes over individual attributes in an entity. The query automatically uses this index when appropriate, which can speed up queries dramatically.

You can specify which attributes to index very simply by using the @Index annotation on the entity attribute:

```

@Entity
public class Customer
{
    @Id String id;
    String firstName;
    @Index String surname;
    String address;
    String phoneNumber;
}

```

The EntityManager creates an appropriate ObjectGrid index for the surname attribute in the Customer entity and the query engine automatically uses the index, which greatly decreases the query time.

Using pagination to improve performance

If there are a million customers named Claus, then it is not likely that you would want to display a page displaying a million customers. It is more likely that you would want to display 10 or 25 customers at a time.

The Query setFirstResult and setMaxResults methods helps by only returning a subset of the results.

Pagination Example

```

Query q = em.createQuery("select c from Customer c where c.surname=:name");
q.setParameter("name", "Claus");
// Display the first page
q.setFirstResult=1;
q.setMaxResults=25;
displayPage(q.getResultIterator());

// Display the second page
q.setFirstResult=26;
displayPage(q.getResultIterator());

```

Reference for eXtreme Scale queries

WebSphere eXtreme Scale has its own language by which the user can query data.

ObjectGrid query FROM clause

The FROM clause specifies the collections of objects to which to apply the query. Each collection is identified either by an abstract schema name and an identification variable, called a range variable, or by a collection member declaration that identifies either a single or multi-valued relationship and an identification variable.

Conceptually, the semantics of the query is to first form a temporary collection of tuples, referred to as R. Tuples are composed of elements from the collections that are identified in the FROM clause. Each tuple contains one element from each of the collections in the FROM clause. All possible combinations are formed subject to the constraints that are imposed by the collection member declarations. If any schema name identifies a collection for which there are no records in the persistent store, then the temporary collection R is empty.

Examples using FROM

The DeptBean object contains records 10, 20 and 30. The EmpBean object contains records 1, 2 and 3 that are related to department 10 and records 4 and 5 that are related to department 20. Department 30 has no related employees.

```
FROM DeptBean d, EmpBean e
```

This clause forms a temporary collection R that contains 15 tuples.

```
FROM DeptBean d, DeptBean d1
```

This clause forms a temporary collection R that contains 9 tuples.

```
FROM DeptBean d, IN (d.emps) AS e
```

This clause forms a temporary collection R that contains 5 tuples. Department 30 is not in the R temporary collection because it contains no employees. Department 10 is contained in the R temporary collection three times and department 20 is contained in R twice.

Instead of using IN(d.emps) as e, you can use a JOIN predicate:

```
FROM DeptBean d JOIN d.emps as e
```

After forming the temporary collection, the search conditions of the WHERE clause are applied to the R temporary collection, yielding a new temporary collection R1. The ORDER BY and SELECT clauses are applied to R1 to yield the final result set.

An identification variable is a variable that is declared in the FROM clause using the IN operator or the optional AS operator.

```
FROM DeptBean AS d, IN (d.emps) AS e
```

is equivalent to:

```
FROM DeptBean d, IN (d.emps) e
```

An identification variable that is declared to be an abstract schema name is called a range variable. In the previous query, "d" is a range variable. An identification variable that is declared to be a multi-valued path expression is called a collection member declaration. The "d" and "e" values in the previous example are collection member declarations.

An example of using a single-valued path expression in the FROM clause follows:

```
FROM EmpBean e, IN(e.dept.mgr) as m
```

ObjectGrid query SELECT clause

The syntax of the SELECT clause is illustrated in the following example:

```
SELECT { ALL | DISTINCT } [ selection , ]* selection
selection ::= {single_valued_path_expression |
               identification_variable |
               OBJECT ( identification_variable) |
               aggregate_functions } [[ AS ] id ]
```

The SELECT clause consists of one or more of the following elements: a single identification variable that is defined in the FROM clause, a single-valued path expression that evaluates to object references or values, and an aggregate function. You can use the DISTINCT keyword to eliminate duplicate references.

A scalar-subselect is a subselect that returns a single value.

Examples using SELECT

Find all employees that earn more than the John employee:

```
SELECT OBJECT(e) FROM EmpBean ej, EmpBean e WHERE ej.name = 'John' and e.salary > ej.salary
```

Find all departments that have one or more employees who earn less than 20000:

```
SELECT DISTINCT e.dept FROM EmpBean e where e.salary < 20000
```

A query can have a path expression that evaluates to an arbitrary value:

```
SELECT e.dept.name FROM EmpBean e where e.salary < 20000
```

The previous query returns a collection of name values for the departments that have employees who earn less than 20000.

A query can return an aggregate value:

```
SELECT avg(e.salary) FROM EmpBean e
```

A query that retrieves the names and object references for underpaid employees follows:

```
SELECT e.name as name, object(e) as emp from EmpBean e where e.salary < 50000
```

ObjectGrid query WHERE clause

The WHERE clause contains search conditions that are composed of the elements presented below. When a search condition evaluates to TRUE, the tuple is added to the result set.

ObjectGrid query literals

A string literal is enclosed in single quotes. A single quotation mark that occurs within a string literal is represented by two single quotes, for example: 'Tom"s'.

A numeric literal can be any of the following values:

- An exact value such as 57, -957, or +66
- Any value supported by Java long type
- A decimal literal such as 57.5 or -47.02
- An approximate numeric value such as 7E3 or -57.4E-2
- Float types must include the "F" qualifier, for example 1.0F
- Long types must include the "L" qualifier, for example 123L

Boolean literals are TRUE and FALSE.

Temporal literals follow JDBC escape syntax based on the type of attribute:

- java.util.Date: yyyy-mm-ss
- java.sql.Date: yyyy-mm-ss
- java.sql.Time: hh-mm-ss
- java.sql.Timestamp: yyyy-mm-dd hh:mm:ss.f...
- java.util.Calendar: yyyy-mm-dd hh:mm:ss.f...

Enum literals are expressed using Java enum literal syntax using the fully qualified enum class name.

ObjectGrid query input parameters

You can specify input parameters by either using an ordinal position or by using a variable name. Writing queries that use input parameters is strongly encouraged, because using input parameters increases performance by allowing the ObjectGrid to catch the query plan between running actions.

An input parameter can be any of the following types: Byte, Short, Integer, Long, Float, Double, BigDecimal, BigInteger, String, Boolean, Char, java.util.Date, java.sql.Date, java.sql.Time, java.sql.Timestamp, java.util.Calendar, a Java SE 5 enum, an Entity or POJO Object, or a binary data string in the form of Java byte[].

An input parameter must not have a NULL value. To search for the occurrence of a NULL value, use the NULL predicate.

Positional Parameters

Positional input parameters are defined by using question mark followed by a positive number:

?[positive integer].

Positional input parameters are numbered starting at 1 and correspond to the arguments of the query; therefore, a query must not contain an input parameter that exceeds the number of input arguments.

Example: `SELECT e FROM Employee e WHERE e.city = ?1 and e.salary >= ?2`

Named Parameters

Named input parameters are defined using a variable name in the format: `:[parameter name]`.

Example: `SELECT e FROM Employee e WHERE e.city = :city and e.salary >= :salary`

ObjectGrid query BETWEEN predicate

The BETWEEN predicate determines whether a given value lies between two other given values.

`expression [NOT] BETWEEN expression-2 AND expression-3`

Example 1

```
e.salary BETWEEN 50000 AND 60000
```

is equivalent to:

```
e.salary >= 50000 AND e.salary <= 60000
```

Example 2

```
e.name NOT BETWEEN 'A' AND 'B'
```

is equivalent to:

```
e.name < 'A' OR e.name > 'B'
```

ObjectGrid query IN predicate

The IN predicate compares a value to a set of values. You can use the IN predicate in one of two forms:

```
expression [NOT] IN ( subselect )expression [NOT] IN ( value1, value2,
.... )
```

The ValueN value can either be a literal value or an input parameter. The expression cannot evaluate to a reference type.

Example 1

```
e.salary IN ( 10000, 15000 )
```

is equivalent to

```
( e.salary = 10000 OR e.salary = 15000 )
```

Example 2

```
e.salary IN ( select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
```

is equivalent to

```
e.salary = ANY ( select e1.salary from EmpBean e1 where e1.dept.deptno =
10)
```

Example 3

```
e.salary NOT IN ( select e1.salary from EmpBean e1 where e1.dept.deptno =
10)
```

is equivalent to

```
e.salary <> ALL ( select e1.salary from EmpBean e1 where e1.dept.deptno =
10)
```

ObjectGrid query LIKE predicate

The LIKE predicate searches a string value for a certain pattern.

```
string-expression [NOT] LIKE pattern [ ESCAPE escape-character ]
```

The pattern value is a string literal or parameter marker of type string in which the underscore (`_`) stands for any single character and percent (`%`) stands for any sequence of characters, including an empty sequence. Any other character stands for itself. The escape character can be used to search for character `_` and `%`. The escape character can be specified as a string literal or as an input parameter.

If the string-expression is null, then the result is unknown.

If both string-expression and pattern are empty, then the result is true.

Example

```
' ' LIKE ' ' is true
' ' LIKE '%' is true
e.name LIKE '12%3' is true for '123' '12993' and false for '1234'
e.name LIKE 's_me' is true for 'some' and 'same', false for 'soome'
e.name LIKE '/_foo' escape '/' is true for '_foo', false for 'afoo'
e.name LIKE '//_foo' escape '/' is true for '/afoo' and for '/bfoo'
e.name LIKE '///_foo' escape '/' is true for '/_foo' but false for '/afoo'
```

ObjectGrid query NULL predicate

The NULL predicate tests for null values.

```
{single-valued-path-expression | input_parameter} IS [NOT] NULL
```

Example

```
e.name IS NULL
e.dept.name IS NOT NULL
e.dept IS NOT NULL
```

ObjectGrid query EMPTY collection predicate

Use the EMPTY collection predicate to test for an empty collection.

To test if a multi-valued relationship is empty, use the following syntax:

```
collection-valued-path-expression IS [NOT] EMPTY
```

Example

Empty collection predicate To find all the departments that have no employees:

```
SELECT OBJECT(d) FROM DeptBean d WHERE d.emps IS EMPTY
```

ObjectGrid query MEMBER OF predicate

The following expression tests whether the object reference that is specified by the single valued path expression or input parameter is a member of the designated collection. If the collection valued path expression designates an empty collection, then the value of the MEMBER OF expression is FALSE.

```
{ single-valued-path-expression | input_parameter } [ NOT ] MEMBER [ OF ]
collection-valued-path-expression
```

Example

Find employees that are not members of a given department number:

```
SELECT OBJECT(e) FROM EmpBean e , DeptBean d
WHERE e NOT MEMBER OF d.emps AND d.deptno = ?1
```

Find employees whose manager is a member of a given department number:

```
SELECT OBJECT(e) FROM EmpBean e, DeptBean d
WHERE e.dept.mgr MEMBER OF d.emps and d.deptno=?1
```

ObjectGrid query EXISTS predicate

The EXISTS predicate tests for the presence or absence of a condition that specified by a subselect.

```
EXISTS ( subselect )
```

The result of EXISTS is true if the subselect returns at least one value, otherwise the result is false.

To negate an EXISTS predicate, precede the predicate with the NOT logical operator.

Example

Return departments that have at least one employee that earns more than 1000000:

```
SELECT OBJECT(d) FROM DeptBean d
WHERE EXISTS ( SELECT e FROM IN (d.emps) e WHERE e.salary > 1000000 )
```

Return departments that have no employees:

```
SELECT OBJECT(d) FROM DeptBean d
WHERE NOT EXISTS ( SELECT e FROM IN (d.emps) e)
```

You can also rewrite the previous query like in the following example:

```
SELECT OBJECT(d) FROM DeptBean d WHERE SIZE(d.emps)=0
```

ObjectGrid query ORDER BY clause

The ORDER BY clause specifies an ordering of the objects in the result collection. An example follows:

```
ORDER BY [ order_element ,]* order_element order_element ::= { path-expression } [
ASC | DESC ]
```

The path expression must specify a single-valued field that is a primitive type of byte, short, int, long, float, double, char, or a wrapper type of Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Character, java.util.Date, java.sql.Date, java.sql.Time, java.sql.Timestamp and java.util.Calendar. The ASC order element specifies that the results are displayed in ascending order, which is the default. A DESC order element specifies that the results are displayed in descending order.

Example

Return department objects. Display the department numbers in decreasing order:

```
SELECT OBJECT(d) FROM DeptBean d ORDER BY d.deptno DESC
```

Return employee objects, sorted by department number and name:

```
SELECT OBJECT(e) FROM EmpBean e ORDER BY e.dept.deptno ASC, e.name DESC
```

ObjectGrid query aggregation functions

Aggregation functions operate on a set of values to return a single scalar value. You can use these functions in the select and subselect methods. The following example illustrates an aggregation:

```
SELECT SUM (e.salary) FROM EmpBean e WHERE e.dept.deptno =20
```

This aggregation computes the total salary for department 20.

The aggregation functions are: AVG, COUNT, MAX, MIN, and SUM. The syntax of an aggregation function is illustrated in the following example:

```
aggregation-function ( [ ALL | DISTINCT ] expression )
```

or:

```
COUNT( [ ALL | DISTINCT ] identification-variable )
```

The DISTINCT option eliminates duplicate values before applying the function. The ALL option is the default option, and does not eliminate duplicate values. Null values are ignored in computing the aggregate function except when you use the COUNT(identification-variable) function, which returns a count of all the elements in the set.

Defining return type

The MAX and MIN functions can apply to any numeric, string or date-time data type and return the corresponding data type. The SUM and AVG functions take a numeric type as input. The AVG function returns a double type. The SUM function returns a long type if the input type is an integer type, except when the input is a Java BigInteger type, then the function returns a Java BigInteger type. The SUM function returns a double type if the input type is not an integer type, except when the input is a Java BigDecimal type, then the function returns a Java BigDecimal type. The COUNT function can take any data type except collections, and returns a long type.

When applied to an empty set, the SUM, AVG, MAX, and MIN functions can return a null value. The COUNT function returns zero (0) when it is applied to an empty set.

Using GROUP BY and HAVING clauses

The set of values that is used for the aggregate function is determined by the collection that results from the FROM and WHERE clause of the query. You can divide the set into groups and apply the aggregation function to each group. To perform this action, use a GROUP BY clause in the query. The GROUP BY clause defines grouping members, which comprise a list of path expressions. Each path expression specifies a field that is a primitive type of byte, short, int, long, float, double, boolean, char, or a wrapper type of Byte, Short, Integer, Long, Float,

Double, BigDecimal, String, Boolean, Character, java.util.Date, java.sql.Date, java.sql.Time, java.sql.Timestamp, java.util.Calendar or a Java SE 5 enum.

The following example illustrates the use of the GROUP BY clause in a query that computes the average salary for each department:

```
SELECT e.dept.deptno, AVG ( e.salary) FROM EmpBean e GROUP BY e.dept.deptno
```

In division of a set into groups, a NULL value is considered equal to another NULL value.

Groups can be filtered using a HAVING clause that tests group properties before involving aggregate functions or grouping members. This filtering is similar to how the WHERE clause filters tuples (that is, records of the return collection values) from the FROM clause. An example of the HAVING clause follows:

```
SELECT e.dept.deptno, AVG ( e.salary) FROM EmpBean e
GROUP BY e.dept.deptno
HAVING COUNT(e) > 3 AND e.dept.deptno > 5
```

This query returns the average salary for departments that have more than three employees and the department number is greater than five.

You can use a HAVING clause without a GROUP BY clause. In this case, the entire set is treated as a single group, to which the HAVING clause is applied.

ObjectGrid query Backus-Naur Form

A summary of the ObjectGrid Query Backus-Naur Form (BNF) Notation follows.

Table 3. Key to BNF summary

| Representation | Description |
|----------------|---------------------|
| {...} | Grouping |
| [...] | Optional constructs |
| bold | Keywords |
| * | Zero or more |
| | Alternates |

```
ObjectGrid QL ::=select_clause from_clause [where_clause] [group_by_clause]
[having_clause] [order_by_clause]
from_clause ::=FROM identification_variable_declaration
[,identification_variable_declaration]*
identification_variable_declaration ::=collection_member_declaration |
range_variable_declaration
collection_member_declaration ::=IN ( collection_valued_path_expression |
single_valued_navigation) [AS] identifier | [LEFT [OUTER]
| INNER] JOIN collection_valued_path_expression |
single_valued_navigation [AS] identifier
range_variable_declaration ::=abstract_schema_name [AS] identifier
single_valued_path_expression ::= {single_valued_navigation | identification_variable}.
{ state_field | state_field.value_object_attribute } | single_valued_navigation
single_valued_navigation ::=identification_variable.[ single_valued_association_field. ]*
single_valued_association_field
collection_valued_path_expression ::=identification_variable.[
single_valued_association_field. ]* collection_valued_association_field
select_clause ::= SELECT [DISTINCT] [ selection , ]* selection
selection ::= {single_valued_path_expression |identification_variable | OBJECT
( identification_variable) |aggregate_functions } [[ AS ] id ]
```

```

order_by_clause ::= ORDER BY [ {identification_variable.[ single_valued_association_field.
]*state_field} [ASC|DESC],]* {identification_variable.[
single_valued_association_field. ]*state_field}[ASC|DESC]
where_clause ::= WHERE conditional_expression
conditional_expression ::= conditional_term | conditional_expression OR conditional_term
conditional_term ::= conditional_factor | conditional_term AND conditional_factor
conditional_factor ::= [NOT] conditional_primary
conditional_primary ::= simple_cond_expression | (conditional_expression)
simple_cond_expression ::= comparison_expression | between_expression | like_expression |
in_expression | null_comparison_expression | empty_collection_comparison_expression |
exists_expression | collection_member_expression
between_expression ::= numeric_expression [NOT] BETWEEN numeric_expression
AND numeric_expression | string_expression [NOT] BETWEEN
string_expression AND string_expression | datetime_expression [NOT]
BETWEEN datetime_expression AND datetime_expression
in_expression ::= identification_variable.[ single_valued_association_field. ]state_field
[*NOT] IN { (subselect) | ( atom ,)* atom }
atom ::= { string_literal | numeric_literal | input_parameter }
like_expression ::=string_expression [NOT] LIKE {string_literal | input_parameter}
[ESCAPE {string_literal | input_parameter}]
null_comparison_expression ::= {single_valued_path_expression | input_parameter} IS
[ NOT ] NULL
empty_collection_comparison_expression ::= collection_valued_path_expression IS
[NOT] EMPTY
collection_member_expression ::= { ssingle_valued_path_expression | input_parameter } [
NOT ] MEMBER [ OF ]collection_valued_path_expression
exists_expression ::= EXISTS {(subselect)}
subselect ::= SELECT [{ ALL | DISTINCT }] subselection from_clause
[where_clause] [group_by_clause] [having_clause]
subselection ::= {single_valued_path_expression |identification_variable |
aggregate_functions }
group_by_clause ::= GROUP BY[single_valued_path_expression,]*
single_valued_path_expression
having_clause ::= HAVING conditional_expression
comparison_expression ::= numeric_expression comparison_operator { numeric_expression
| {SOME | ANY | ALL} (subselect) } | string_expression
comparison_operator {
string_expression | {SOME | ANY | ALL}(subselect) } |
datetime_expression comparison_operator {
datetime_expression {SOME | ANY | ALL}(subselect) } |
boolean_expression {=|<>} {
boolean_expression {SOME | ANY | ALL}(subselect) } |
entity_expression {=|<>} {
entity_expression {SOME| ANY | ALL}(subselect) }
comparison_operator ::= = | > | >= | < | <= | <>
string_expression ::= string_primary | (subselect)
string_primary ::=state_field_path_expression |string_literal | input_parameter |
functions_returning_strings
datetime_expression ::= datetime_primary |(subselect)
datetime_primary ::=state_field_path_expression | string_literal | long_literal
| input_parameter | functions_returning_datetime
boolean_expression ::= boolean_primary |(subselect)
boolean_primary ::=state_field_path_expression | boolean_literal | input_parameter
entity_expression ::=single_valued_association_path_expression |
identification_variable | input_parameter
numeric_expression ::= simple_numeric_expression |(subselect)
simple_numeric_expression ::= numeric_term | numeric_expression {+|-} numeric_term
numeric_term ::= numeric_factor | numeric_term {*/|} numeric_factor
numeric_factor ::= {+|-} numeric_primary

```

```

numeric_primary ::= single_valued_path_expression | numeric_literal |
  ( numeric_expression ) | input_parameter | functions
aggregate_functions ::=
AVG([ALL|DISTINCT] identification_variable.
  [ single_valued_association_field. ]*state_field) |
COUNT([ALL|DISTINCT] {single_valued_path_expression |
  identification_variable}) |
MAX([ALL|DISTINCT] identification_variable.[
  single_valued_association_field. ]*state_field) |
MIN([ALL|DISTINCT] identification_variable.[
  single_valued_association_field. ]*state_field) |
SUM([ALL|DISTINCT] identification_variable.[
  single_valued_association_field. ]*state_field)
functions ::=
ABS (simple_numeric_expression) |
CONCAT (string_primary , string_primary) |
LOWER (string_primary) |
LENGTH(string_primary) |
LOCATE(string_primary, string_primary [, simple_numeric_expression]) |
MOD (simple_numeric_expression, simple_numeric_expression) |
SIZE (collection_valued_path_expression) |
SQRT (simple_numeric_expression) |
SUBSTRING (string_primary, simple_numeric_expression[, simple_numeric_expression]) |
UPPER (string_primary) |
TRIM ([[LEADING | TRAILING | BOTH] [trim_character]
FROM] string_primary)

```

Query performance tuning

To tune the performance of your queries, use the following techniques and tips.

Using parameters

When a query runs, the query string must be parsed and a plan developed to run the query, both of which can be costly. WebSphere eXtreme Scale caches query plans by the query string. Since the cache is a finite size, it is important to reuse query strings whenever possible. Using named or positional parameters also helps performance by fostering query plan reuse.

Positional Parameter Example Query `q = em.createQuery("select c from Customer c where c.surname=?1"); q.setParameter(1, "Claus");`

Using indexes

Proper indexing on a map might have a significant impact on query performance, even though indexing has some overhead on overall map performance. Without indexing on object attributes involved in queries, the query engine performs a table scan for each attribute. The table scan is the most expensive operation during a query run. Indexing on object attributes that are involved in queries allow the query engine to avoid an unnecessary table scan, improving the overall query performance. If the application is designed to use query intensively on a read-most map, configure indexes for object attributes that are involved in the query. If the map is mostly updated, then you must balance between query performance improvement and indexing overhead on the map. See Indexing for more information.

When plain old Java objects (POJO) are stored in a map, proper indexing can avoid a Java reflection. In the following example, query replaces the WHERE clause with

range index search, if the budget field has an index built over it. Otherwise, query scans the entire map and evaluates the WHERE clause by first getting the budget using Java reflection and then comparing the budget with the value 50000:

```
SELECT d FROM DeptBean d WHERE d.budget=50000
```

See “Query plan” for details on how to best tune individual queries and how different syntax, object models and indexes can affect query performance.

Using pagination

In client-server environments, the query engine transports the entire result map to the client. The data that is returned should be divided into reasonable chunks. The EntityManager Query and ObjectMap ObjectQuery interfaces both support the setFirstResult and setMaxResults methods that allow the query to return a subset of the results.

Return primitive values instead of entities

With the EntityManager Query API, entities are returned as query parameters. The query engine currently returns the keys for these entities to the client. When the client iterates over these entities using the Iterator from the getResultIterator method, each entity is automatically inflated and managed as if it were created with the find method on the EntityManager interface. The entire entity graph is built from the entity ObjectMap on the client. The entity value attributes and any related entities are eagerly resolved.

To avoid building the costly graph, modify the query to return the individual attributes with path navigation.

For example:

```
// Returns an entity
SELECT p FROM Person p
// Returns attributes SELECT p.name, p.address.street, p.address.city, p.gender FROM Person p
```

Query plan

All eXtreme Scale queries have a query plan. The plan describes how the query engine interacts with ObjectMaps and indexes. Display the query plan to determine if the query string or indexes are being used appropriately. The query plan can also be used to explore the differences that subtle changes in a query string make in the way eXtreme Scale runs a query.

The query plan can be viewed one of two ways:

- EntityManager Query or ObjectQuery getPlan API methods
- ObjectGrid diagnostic trace

getPlan method

The getPlan method on the ObjectQuery and Query interfaces return a String that describes the query plan. This string can be displayed to standard output or a log to display a query plan.

Note: In a distributed environment, the getPlan method does not run against the server and does not reflect any defined indexes. To view the plan, use an agent to view the plan on the server.

Query plan trace

The query plan can be displayed using ObjectGrid trace. To enable query plan trace, use the following trace specification:

```
QueryEnginePlan=debug=enabled
```

See “Collecting trace” on page 358 for details on how to enable trace and locate the trace log files.

Query plan examples

Query plan uses the word `for` to indicate that the query is iterating through an `ObjectMap` collection or through a derived collection such as `q2.getEmps()`, `q2.dept`, or a temporary collection returned by an inner loop. If the collection is from an `ObjectMap`, the query plan shows whether a sequential scan (denoted by `INDEX SCAN`), unique or non-unique index is used. Query plan uses a filter string to list the condition expressions applied to a collection.

A Cartesian product is not commonly used in object query. The following query scans the entire `EmpBean` map in the outer loop and scans the entire `DeptBean` map in the inner loop:

```
SELECT e, d FROM EmpBean e, DeptBean d
```

Plan trace:

```
for q2 in EmpBean ObjectMap using INDEX SCAN
  for q3 in DeptBean ObjectMap using INDEX SCAN
    returning new Tuple( q2, q3 )
```

The following query retrieves all employee names from a particular department by sequentially scanning the `EmpBean` map to get an employee object. From the employee object, the query navigates to its department object and applies the `d.no=1` filter. In this example, each employee has only one department object reference, so the inner loop runs one time:

```
SELECT e.name FROM EmpBean e JOIN e.dept d WHERE d.no=1
```

Plan trace:

```
for q2 in EmpBean ObjectMap using INDEX SCAN
  for q3 in q2.dept
    filter ( q3.getNo() = 1 )
    returning new Tuple( q2.name )
```

The following query is equivalent to the previous query. However, the following query performs better because it first narrows the result down to one department object by using the unique index that is defined over the `DeptBean` primary key field number. From the department object, the query navigates to its employee objects to get their names:

```
SELECT e.name FROM DeptBean d JOIN d.emps e WHERE d.no=1
```

Plan trace:

```
for q2 in DeptBean ObjectMap using UNIQUE INDEX key=(1)
  for q3 in q2.getEmps()
    returning new Tuple( q3.name )
```


The following query finds all the employees that work for development or sales. The query scans the entire EmpBean map and performs additional filtering by evaluating the expressions: d.name = 'Sales' or d.name='Dev'

```
SELECT e FROM EmpBean e, in (e.dept) d WHERE d.name = 'Sales'
or d.name='Dev'
```

Plan trace:

```
for q2 in EmpBean ObjectMap using INDEX SCAN
  for q3 in q2.dept
    filter (( q3.getName() = Sales ) OR ( q3.getName() = Dev ) )
  returning new Tuple( q2 )
```

The following query is equivalent to the previous query, but this query runs a different query plan and uses the range index built over the field name. In general, this query performs better because the index over the name field is used for narrowing down the department objects, which run quickly if only a few departments are development or sales.

```
SELECT e FROM DeptBean d, in(d.emps) e WHERE d.name='Dev' or d.name='Sales'
```

Plan trace:

IteratorUnionIndex of

```
    for q2 in DeptBean ObjectMap using INDEX on name = (Dev)
      for q3 in q2.getEmps()

    for q2 in DeptBean ObjectMap using INDEX on name = (Sales)
      for q3 in q2.getEmps()
```

The following query finds departments that do not have any employees:

```
SELECT d FROM DeptBean d WHERE NOT EXISTS(select e from d.emps e)
```

Plan trace:

```
for q2 in DeptBean ObjectMap using INDEX SCAN
  filter ( NOT EXISTS ( correlated collection defined as

    for q3 in q2.getEmps()
    returning new Tuple( q3 )

  returning new Tuple( q2 )
```

The following query is equivalent to the previous query but uses the SIZE scalar function. This query has similar performance but is easier to write.

```
SELECT d FROM DeptBean d WHERE SIZE(d.emps)=0
for q2 in DeptBean ObjectMap using INDEX SCAN
  filter (SIZE( q2.getEmps()) = 0 )
  returning new Tuple( q2 )
```

The following example is another way of writing the same query as the previous query with similar performance, but this query is easier to write as well:

```
SELECT d FROM DeptBean d WHERE d.emps is EMPTY
```

Plan trace:

```
for q2 in DeptBean ObjectMap using INDEX SCAN
  filter ( q2.getEmps() IS EMPTY )
  returning new Tuple( q2 )
```

The following query finds any employees with a home address matching at least one of the addresses of the employee whose name equals the value of the parameter. The inner loop has no dependency on the outer loop. The query runs the inner loop one time.

```
SELECT e FROM EmpBean e WHERE e.home = any (SELECT e1.home FROM EmpBean e1
WHERE e1.name=?1)
for q2 in EmpBean ObjectMap using INDEX SCAN
  filter ( q2.home =ANY      temp collection defined as

      for q3 in EmpBean ObjectMap using INDEX on name = ( ?1)
      returning new Tuple( q3.home      )
  )
returning new Tuple( q2 )
```

The following query is equivalent to the previous query, but has a correlated subquery; also, the inner loop runs repeatedly.

```
SELECT e FROM EmpBean e WHERE EXISTS(SELECT e1 FROM EmpBean e1 WHERE
e.home=e1.home and e1.name=?1)
```

Plan trace:

```
for q2 in EmpBean ObjectMap using INDEX SCAN
  filter ( EXISTS (      correlated collection defined as

      for q3 in EmpBean ObjectMap using INDEX on name = (?1)
      filter ( q2.home = q3.home )
      returning new Tuple( q3      )

  )
returning new Tuple( q2 )
```

Query optimization using indexes

Defining and using indexes properly can significantly improve query performance.

WebSphere eXtreme Scale queries can use built-in HashIndex plug-ins to improve performance of queries. Indexes can be defined on entity or object attributes. The query engine will automatically use the defined indexes if its WHERE clause uses one of the following strings:

- A comparison expression with the following operators: =, <, >, <= or >= (any comparison expressions except not equals <>)
- A BETWEEN expression
- Operands of the expressions are constants or simple terms

Requirements

Indexes have the following requirements when used by Query:

- All indexes must use the built-in HashIndex plug-in.
- All indexes must be statically defined. Dynamic indexes are not supported.
- The @Index annotation may be used to automatically create static HashIndex plug-ins.
- All single-attribute indexes must have the RangeIndex property set to true.
- All composite indexes must have the RangeIndex property set to false.
- All association (relationship) indexes must have the RangeIndex property set to false.

For information about configuring the HashIndex, refer to Configuring the HashIndex.

For information regarding indexing, see Indexing.

For a more efficient way to search for cached objects, see “Using a composite index” on page 227

Using hints to choose an index

An index can be manually selected using the `setHint` method on the `Query` and `ObjectQuery` interfaces with the `HINT_USEINDEX` constant. This can be helpful when optimizing a query to use the best performing index.

Query examples that use attribute indexes

The following examples use simple terms: `e.empid`, `e.name`, `e.salary`, `d.name`, `d.budget` and `e.isManager`. The examples assume that indexes are defined over the `name`, `salary` and `budget` fields of an entity or value object. The `empid` field is a primary key and `isManager` has no index defined.

The following query uses both indexes over the fields of `name` and `salary`. It returns all employees with names that equal the value of the first parameter or a salary equal to the value of the second parameter:

```
SELECT e FROM EmpBean e where e.name=?1 or e.salary=?2
```

The following query uses both indexes over the fields of `name` and `budget`. The query returns all departments named 'DEV' with a budget that is greater than 2000.

```
SELECT d FROM DeptBean dwhere d.name='DEV' and d.budget>2000
```

The following query returns all employees with a salary greater than 3000 and with an `isManager` flag value that equals the value of the parameter. The query uses the index that is defined over the `salary` field and performs additional filtering by evaluating the comparison expression: `e.isManager=?1`.

```
SELECT e FROM EmpBean e where e.salary>3000 and e.isManager=?1
```

The following query finds all employees who earn more than the first parameter, or any employee that is a manager. Although the `salary` field has an index defined, query scans the built-in index that is built over the primary keys of the `EmpBean` field and evaluates the expression: `e.salary>?1 or e.isManager=TRUE`.

```
SELECT e FROM EmpBean e WHERE e.salary>?1 or e.isManager=TRUE
```

The following query returns employees with a name that contains the letter `a`. Although the `name` field has an index defined, query does not use the index because the `name` field is used in the `LIKE` expression.

```
SELECT e FROM EmpBean e WHERE e.name LIKE '%a%'
```

The following query finds all employees with a name that is not "Smith". Although the `name` field has an index defined, query does not use the index because the query uses the not equals (`<>`) comparison operator.

```
SELECT e FROM EmpBean e where e.name<>'Smith'
```

The following query finds all departments with a budget less than the value of the parameter, and with an employee salary greater than 3000. The query uses an index for the salary, but it does not use an index for the budget because dept.budget is not a simple term. The dept objects are derived from collection e. You do not need to use the budget index to look for dept objects.

```
SELECT dept from EmpBean e, in (e.dept) dept where e.salary>3000 and
dept.budget<?
```

The following query finds all employees with a salary greater than the salary of the employees that have the empid of 1, 2, and 3. The index salary is not used because the comparison involves a subquery. The empid is a primary key, however, and is used for a unique index search because all the primary keys have a built-in index defined.

```
SELECT e FROM EmpBean e WHERE e.salary > ALL (SELECT e1.salary FROM
EmpBean e1 WHERE e1.empid=1 or e1.empid =2 or e1.empid=99)
```

To check if the index is being used by the query, you can view the “Query plan” on page 117. Here is an example query plan for the previous query:

```
for q2 in EmpBean ObjectMap using INDEX SCAN
  filter ( q2.salary >ALL temp collection defined as
    IteratorUnionIndex of
      for q3 in EmpBean ObjectMap using UNIQUE INDEX key=(1)
      )
      for q3 in EmpBean ObjectMap using UNIQUE INDEX key=(2)
      )
      for q3 in EmpBean ObjectMap using UNIQUE INDEX key=(99)
      )
  returning new Tuple( q3.salary )
returning new Tuple( q2 )

for q2 in EmpBean ObjectMap using RANGE INDEX on salary with range(3000,)
  for q3 in q2.dept
    filter ( q3.budget < ?1 )
  returning new Tuple( q3 )
```

Indexing attributes

Indexes can be defined over any single attribute type with the constraints previously defined.

Defining entity indexes using @Index

To define an index on an entity, simply define an annotation:

Entities using annotations

```
@Entity
public class Employee {
  @Id int empid;
  @Index String name
  @Index double salary
  @ManyToOne Department dept;
}
@Entity
public class Department {
  @Id int deptid;
  @Index String name;
```

```

@Index double budget;
boolean isManager;
@OneToMany Collection<Employee> employees;
}

```

With XML

Indexes can also be defined using XML:

Entities without annotations

```

public class Employee {
int empid;
String name;
double salary;
Department dept;
}

```

```

public class Department {
int deptid;
String name;
double budget;
boolean isManager;
Collection employees;
}

```

ObjectGrid XML with attribute indexes

```

<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
xmlns="http://ibm.com/ws/objectgrid/config">
<objectGrids>
<objectGrid name="DepartmentGrid" entityMetadataXMLFile="entity.xml">
<backingMap name="Employee" pluginCollectionRef="Emp"/>
<backingMap name="Department" pluginCollectionRef="Dept"/>
</objectGrid>
</objectGrids>
<backingMapPluginCollections>
<backingMapPluginCollection id="Emp">
<bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
<property name="Name" type="java.lang.String" value="Employee.name"/>
<property name="AttributeName" type="java.lang.String" value="name"/>
<property name="RangeIndex" type="boolean" value="true"
description="Ranges are must be set to true for attributes." />
</bean>
<bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
<property name="Name" type="java.lang.String" value="Employee.salary"/>
<property name="AttributeName" type="java.lang.String" value="salary"/>
<property name="RangeIndex" type="boolean" value="true"
description="Ranges are must be set to true for attributes." />
</bean>
</backingMapPluginCollection>
<backingMapPluginCollection id="Dept">
<bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
<property name="Name" type="java.lang.String" value="Department.name"/>
<property name="AttributeName" type="java.lang.String" value="name"/>
<property name="RangeIndex" type="boolean" value="true"
description="Ranges are must be set to true for attributes." />
</bean>
<bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
<property name="Name" type="java.lang.String" value="Department.budget"/>
<property name="AttributeName" type="java.lang.String" value="budget"/>
<property name="RangeIndex" type="boolean" value="true"
description="Ranges are must be set to true for attributes." />
</bean>
</backingMapPluginCollection>
</backingMapPluginCollections>
</objectGridConfig>

```

Entity XML

```

<?xml version="1.0" encoding="UTF-8"?>
<entity-mappings xmlns="http://ibm.com/ws/projector/config/emd"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/ws/projector/config/emd ../emd.xsd">
<description>Department entities</description>
<entity class-name="acme.Employee" name="Employee" access="FIELD">
<attributes>
<id name="empid" />
<basic name="name" />
<basic name="salary" />

```

```

<many-to-one name="department"
target-entity="acme.Department"
fetch="EAGER">
<cascade><cascade-persist/></cascade>
</many-to-one>
</attributes>
</entity>
<entity class-name="acme.Department" name="Department" access="FIELD">
<attributes>
<id name="deptid" />
<basic name="name" />
<basic name="budget" />
<basic name="isManager" />
<one-to-many name="employees"
target-entity="acme.Employee"
fetch="LAZY" mapped-by="parentNode">
<cascade><cascade-persist/></cascade>
</one-to-many>
</attributes>
</entity>
</entity-mappings>

```

Defining indexes for non-entities using XML

Indexes for non-entity types are defined in XML. There is no difference when creating the MapIndexPlugin for entity maps and non-entity maps.

Java bean

```

public class Employee {
    int empid;
    String name;
    double salary;
    Department dept;

    public class Department {
        int deptid;
        String name;
        double budget;
        boolean isManager;
        Collection employees;
    }
}

```

ObjectGrid XML with attribute indexes

```

<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
xmlns="http://ibm.com/ws/objectgrid/config">
<objectGrids>
<objectGrid name="DepartmentGrid">
<backingMap name="Employee" pluginCollectionRef="Emp"/>
<backingMap name="Department" pluginCollectionRef="Dept"/>
<querySchema>
<mapSchemas>
<mapSchema mapName="Employee" valueClass="acme.Employee"
primaryKeyField="empid" />
<mapSchema mapName="Department" valueClass="acme.Department"
primaryKeyField="deptid" />
</mapSchemas>
<relationships>
<relationship source="acme.Employee"
target="acme.Department"
relationField="dept" invRelationField="employees" />
</relationships>
</querySchema>
</objectGrid>
</objectGrids>
<backingMapPluginCollections>
<backingMapPluginCollection id="Emp">
<bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
<property name="Name" type="java.lang.String" value="Employee.name"/>
<property name="AttributeName" type="java.lang.String" value="name"/>
<property name="RangeIndex" type="boolean" value="true"
description="Ranges are must be set to true for attributes." />
</bean>
<bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
<property name="Name" type="java.lang.String" value="Employee.salary"/>
<property name="AttributeName" type="java.lang.String" value="salary"/>
<property name="RangeIndex" type="boolean" value="true"
description="Ranges are must be set to true for attributes." />
</bean>
</backingMapPluginCollection>
<backingMapPluginCollection id="Dept">
<bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">

```

```

<property name="Name" type="java.lang.String" value="Department.name"/>
<property name="AttributeName" type="java.lang.String" value="name"/>
<property name="RangeIndex" type="boolean" value="true"
description="Ranges are must be set to true for attributes." />
</bean>
<bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
<property name="Name" type="java.lang.String" value="Department.budget"/>
<property name="AttributeName" type="java.lang.String" value="budget"/>
<property name="RangeIndex" type="boolean" value="true"
description="Ranges are must be set to true for attributes." />
</bean>
</backingMapPluginCollection>
</backingMapPluginCollections>
</objectGridConfig>

```

Indexing relationships

WebSphere eXtreme Scale stores the foreign keys for related entities within the parent object. For entities, the keys are stored in the underlying tuple. For non-entity objects, the keys are explicitly stored in the parent object.

Adding an index on a relationship attribute can speed up queries that use cyclical references or use the IS NULL, IS EMPTY, SIZE and MEMBER OF query filters. Both single- and multi-valued associations may have the @Index annotation or a HashIndex plug-in configuration in an ObjectGrid descriptor XML file.

Defining entity relationship indexes using @Index

The following example defines entities with @Index annotations:

Entity with annotation

```

@Entity
public class Node {
    @ManyToOne @Index
    Node parentNode;

    @OneToMany @Index
    List<Node> childrenNodes = new ArrayList();

    @OneToMany @Index
    List<BusinessUnitType> businessUnitTypes = new ArrayList();
}

```

Defining entity relationship indexes using XML

The following example defines the same entities and indexes using XML with HashIndex plug-ins:

Entity without annotations

```

public class Node {
    int nodeId;
    Node parentNode;
    List<Node> childrenNodes = new ArrayList();
    List<BusinessUnitType> businessUnitTypes = new ArrayList();
}

```

ObjectGrid XML

```

<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
xmlns="http://ibm.com/ws/objectgrid/config">
<objectGrids>
<objectGrid name="ObjectGrid_Entity" entityMetadataXMLFile="entity.xml">
<backingMap name="Node" pluginCollectionRef="Node"/>
<backingMap name="BusinessUnitType" pluginCollectionRef="BusinessUnitType"/>
</objectGrid>
</objectGrids>
<backingMapPluginCollections>
<backingMapPluginCollection id="Node">

```

```

<bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
<property name="Name" type="java.lang.String" value="parentNode"/>
<property name="AttributeName" type="java.lang.String" value="parentNode"/>
</property name="RangeIndex" type="boolean" value="false"
description="Ranges are not supported for association indexes." /> </bean>
<bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
<property name="Name" type="java.lang.String" value="businessUnitType"/>
<property name="AttributeName" type="java.lang.String" value="businessUnitTypes"/>
</property name="RangeIndex" type="boolean" value="false"
description="Ranges are not supported for association indexes." />
</bean>
</backingMapPluginCollection>
</backingMapPluginCollections>
</objectGridConfig>

```

Entity XML

```

<?xml version="1.0" encoding="UTF-8"?>
<entity-mappings xmlns="http://ibm.com/ws/projector/config/emd"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/ws/projector/config/emd ./emd.xsd">
<description>My entities</description>
<entity class-name="acme.Node" name="Account" access="FIELD">
<attributes>
<id name="nodeId" />
<one-to-many name="childrenNodes"
target-entity="acme.Node"
fetch="EAGER" mapped-by="parentNode">
<cascade><cascade-all/></cascade>
</one-to-many>
<many-to-one name="parentNodes"
target-entity="acme.Node"
fetch="LAZY" mapped-by="childrenNodes">
<cascade><cascade-none/></cascade>
</many-to-one>
<many-to-one name="businessUnitTypes"
target-entity="acme.BusinessUnitType"
fetch="EAGER">
<cascade><cascade-persist/></cascade>
</many-to-one>
</attributes>
</entity>
<entity class-name="acme.BusinessUnitType" name="BusinessUnitType" access="FIELD">
<attributes>
<id name="build" />
<basic name="TypeDescription" />
</attributes>
</entity>
</entity-mappings>

```

Using the previously defined indexes, the following entity query examples are optimized:

```

SELECT n FROM Node n WHERE n.parentNode is null
SELECT n FROM Node n WHERE n.businessUnitTypes is EMPTY
SELECT n FROM Node n WHERE size(n.businessUnitTypes)>=10
SELECT n FROM BusinessUnitType b, Node n WHERE b member of n.businessUnitTypes and b.name='TELECOM'

```

Defining non-entity relationship indexes

The following example defines a HashIndex plug-in for non-entity maps in an ObjectGrid descriptor XML file:

```

<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
xmlns="http://ibm.com/ws/objectgrid/config">
<objectGrids>
<objectGrid name="ObjectGrid_POJO">
<backingMap name="Node" pluginCollectionRef="Node"/>
<backingMap name="BusinessUnitType" pluginCollectionRef="BusinessUnitType"/>
<querySchema>
<mapSchemas>
<mapSchema mapName="Node"
valueClass="com.ibm.websphere.objectgrid.samples.entity.Node"
primaryKeyField="id" />
<mapSchema mapName="BusinessUnitType"
valueClass="com.ibm.websphere.objectgrid.samples.entity.BusinessUnitType"
primaryKeyField="id" />

```



```

</mapSchemas>
<relationships>
  <relationship source="com.ibm.websphere.objectgrid.samples.entity.Node"
    target="com.ibm.websphere.objectgrid.samples.entity.Node"
    relationField="parentNodeId" invRelationField="childrenNodeIds" />
  <relationship source="com.ibm.websphere.objectgrid.samples.entity.Node"
    target="com.ibm.websphere.objectgrid.samples.entity.BusinessUnitType"
    relationField="businessUnitTypeKeys" invRelationField="" />
</relationships>
</querySchema>
</objectGrid>
</objectGrids>
<backingMapPluginCollections>
  <backingMapPluginCollection id="Node">
    <bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
      <property name="Name" type="java.lang.String" value="parentNode"/>
    </bean>
    <property name="Name" type="java.lang.String" value="parentNodeId"/>
    <property name="AttributeName" type="java.lang.String" value="parentNodeId"/>
    <property name="RangeIndex" type="boolean" value="false"
      description="Ranges are not supported for association indexes." />
  </bean>
  <bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
    <property name="Name" type="java.lang.String" value="businessUnitType"/>
    <property name="AttributeName" type="java.lang.String" value="businessUnitTypeKeys"/>
  </bean>
  <property name="RangeIndex" type="boolean" value="false"
    description="Ranges are not supported for association indexes." />
  </bean>
  <bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
    <property name="Name" type="java.lang.String" value="childrenNodeIds"/>
    <property name="AttributeName" type="java.lang.String" value="childrenNodeIds"/>
    <property name="RangeIndex" type="boolean" value="false"
      description="Ranges are not supported for association indexes." />
  </bean>
</backingMapPluginCollection>
</backingMapPluginCollections>
</objectGridConfig>

```

Given the above index configurations, the following object query examples are optimized:

```

SELECT n FROM Node n WHERE n.parentNodeId is null
SELECT n FROM Node n WHERE n.businessUnitTypeKeys is EMPTY
SELECT n FROM Node n WHERE size(n.businessUnitTypeKeys)>=10
SELECT n FROM BusinessUnitType b, Node n WHERE
  b member of n.businessUnitTypeKeys and b.name='TELECOM'

```

Using objects other than keys to find partitions (PartitionableKey interface)

When an eXtreme Scale configuration uses the fixed partition placement strategy, it depends on hashing the key to a partition to insert, get, update, or remove the value. The hashCode method is called on the key and it must be well defined if a custom key is created. However another option is to use the PartitionableKey interface. With the PartitionableKey interface, you can use an object other than the key to hash to a partition.

You can use the PartitionableKey interface in situations where there are multiple maps and the data you commit is related and thus should be put on the same partition. WebSphere eXtreme Scale does not support two-phase commit so multiple map transactions should not be committed if they span multiple partitions. If the PartitionableKey hashes to the same partition for keys in different maps in the same map set, they can be committed together.

You can also use the PartitionableKey interface when groups of keys should be put on the same partition, but not necessarily during a single transaction. If keys should be hashed based on location, department, domain type, or some other type of identifier, children keys can be given a parent PartitionableKey.

For example, employees should hash to the same partition as their department. Each employee key would have a `PartitionableKey` object that belongs to the department map. Then both the employee and department would hash to the same partition.

The `PartitionableKey` interface supplies one method, called `ibmGetPartition`. The object returned from this method must implement the `hashCode` method. The result returned from using the alternate `hashCode` will be used to route the key to a partition.

Programming for transactions

Applications that require transactions introduce such considerations as handling locks, handling collisions, and transaction isolation.

Transaction processing overview

WebSphere eXtreme Scale uses transactions as its mechanism for interaction with data.

To interact with data, the thread in your application needs its own session. When the application wants to use the `ObjectGrid` on a thread, call one of the `ObjectGrid.getSession` methods to obtain a thread. With the session, the application can work with data that is stored in the `ObjectGrid` maps.

When an application uses a `Session` object, the session must be in the context of a transaction. A transaction begins and commits or begins and rolls back using the `begin`, `commit`, and `rollback` methods on the `Session` object. Applications can also work in auto-commit mode, in which the `Session` automatically begins and commits a transaction whenever an operation is performed on the map. Auto-commit mode cannot group multiple operations into a single transaction, so it is the slower option if you are creating a batch of multiple operations into a single transaction. However, for transactions that contain only one operation, auto-commit is the faster option.

Transactions

Transactions have many advantages for data storage and manipulation. You can use transactions to protect the data grid from concurrent changes, to apply multiple changes as a concurrent unit, to replicate data and to implement a life cycle for locks on changes.

When a transaction starts, WebSphere eXtreme Scale allocates a special difference map to hold the current changes or copies of key and value pairs that the transaction uses. Typically, when a key and value pair is accessed, the value is copied before the application receives the value. The difference map tracks all changes for operations such as `insert`, `update`, `get`, `remove`, and so on. Keys are not copied because they are assumed to be immutable. If an `ObjectTransformer` object is specified, then this object is used for copying the value. If the transaction is using optimistic locking, then before images of the values are also tracked for comparison when the transaction commits.

If a transaction is rolled back, then the difference map information is discarded, and locks on entries are released. When a transaction commits, the changes are applied to the maps and locks are released. If optimistic locking is being used, then eXtreme Scale compares the before image versions of the values with the values that are in the map. These values must match for the transaction to commit. This

comparison enables a multiple version locking scheme, but at a cost of two copies being made when the transaction accesses the entry. All values are copied again and the new copy is stored in the map. WebSphere eXtreme Scale performs this copy to protect itself against the application changing the application reference to the value after a commit.

You can avoid using several copies of the information. The application can save a copy by using pessimistic locking instead of optimistic locking as the cost of limiting concurrency. The copy of the value at commit time can also be avoided if the application agrees not to change a value after a commit.

Advantages of transactions

Use transactions for the following reasons:

By using transactions, you can:

- Roll back changes if an exception occurs or business logic needs to undo state changes.
- To apply multiple changes as an atomic unit at commit time.
- Hold and release locks on data to apply multiple changes as an atomic unit at commit time.
- Protect a thread from concurrent changes.
- Implement a life cycle for locks on changes.
- Produce an atomic unit of replication.

Transaction size

Larger transactions are more efficient, especially for replication. However, larger transactions can adversely impact concurrency because the locks on entries are held for a longer period of time. If you use larger transactions, you can increase replication performance. This performance increase is important when you are pre-loading a Map. Experiment with different batch sizes to determine what works best for your scenario.

Larger transactions also help with loaders. If a loader is being used that can perform SQL batching, then significant performance gains are possible depending on the transaction and significant load reductions on the database side. This performance gain depends on the Loader implementation.

Automatic commit mode

If no transaction is actively started, then when an application interacts with an ObjectMap object, an automatic begin and commit operation is done on behalf of the application. This automatic begin and commit operation works, but prevents rollback and locking from working effectively. Synchronous replication speed is impacted because of the very small transaction size. If you are using an entity manager application, then do not use automatic commit mode because objects that are looked up with the EntityManager.find method immediately become unmanaged on the method return and become unusable.

External transaction coordinators

Typically, transactions begin with the session.begin method and end with the session.commit method. However, when eXtreme Scale is embedded, the

transactions might be started and ended by an external transaction coordinator. If you are using an external transaction coordinator, you do not need to call the `session.begin` method and end with the `session.commit` method. If you are using WebSphere Application Server, you can use the `WebSphereTransactionCallback` plug-in.

CopyMode attribute

You can tune the number of copies by defining the `CopyMode` attribute of the `BackingMap` or `ObjectMap` objects.

You can tune the number of copies by defining the `CopyMode` attribute of the `BackingMap` or `ObjectMap` objects. The copy mode has the following values:

- `COPY_ON_READ_AND_COMMIT`
- `COPY_ON_READ`
- `NO_COPY`
- `COPY_ON_WRITE`
- `COPY_TO_BYTES`

The `COPY_ON_READ_AND_COMMIT` value is the default. The `COPY_ON_READ` value copies on the initial data retrieved, but does not copy at commit time. This mode is safe if the application does not modify a value after committing a transaction. The `NO_COPY` value does not copy data, which is only safe for read-only data. If the data never changes then you do not need to copy it for isolation reasons.

Be careful when you use the `NO_COPY` attribute value with maps that can be updated. WebSphere eXtreme Scale uses the copy on first touch to allow the transaction rollback. The application only changed the copy, and as a result, eXtreme Scale discards the copy. If the `NO_COPY` attribute value is used, and the application modifies the committed value, completing a rollback is not possible. Modifying the committed value leads to problems with indexes, replication, and so on because the indexes and replicas update when the transaction commits. If you modify committed data and then roll back the transaction, which does not actually roll back at all, then the indexes are not updated and replication does not take place. Other threads can see the uncommitted changes immediately, even if they have locks. Use the `NO_COPY` attribute value for read-only maps or for applications that complete the appropriate copy before modifying the value. If you use the `NO_COPY` attribute value and call IBM support with a data integrity problem, you are asked to reproduce the problem with the copy mode set to `COPY_ON_READ_AND_COMMIT`.

The `COPY_TO_BYTES` value stores values in the map in a serialized form. At read time, eXtreme Scale inflates the value from a serialized form and at commit time it stores the value to a serialized form. With this method, a copy occurs at both read and commit time.

The default copy mode for a map can be configured on the `BackingMap` object. You can also change the copy mode on maps before you start a transaction by using the `ObjectMap.setCopyMode` method.

An example of a backing map snippet from an `objectgrid.xml` file that shows how to set the copy mode for a given backing map follows. This example assumes that you are using `cc` as the `objectgrid/config` namespace.

```
<cc:backingMap name="RuntimeLifespan" copyMode="NO_COPY"/>
```

See the information about copyMode best practices in the *Programming Guide* for more information.

Map entry locking

An ObjectGrid BackingMap supports several locking strategies for maps to maintain cache entry consistency.

Lock manager configuration

When either a PESSIMISTIC or an OPTIMISTIC lock strategy is used, a lock manager is created for the BackingMap. The lock manager uses a hash map to track entries that are locked by one or more transactions. If many map entries exist in the hash map, more lock buckets can result in better performance. The risk of Java synchronization collisions is lower as the number of buckets grows. More lock buckets also lead to more concurrency. The previous examples show how an application can set the number of lock buckets to use for a given BackingMap instance.

To avoid a java.lang.IllegalStateException exception, the setNumberOfLockBuckets method must be called before calling the initialize or getSession methods on the ObjectGrid instance. The setNumberOfLockBuckets method parameter is a Java primitive integer that specifies the number of lock buckets to use. Using a prime number can allow for a uniform distribution of map entries over the lock buckets. A good starting point for best performance is to set the number of lock buckets to about 10 percent of the expected number of BackingMap entries.

Locking strategies

Locking strategies include pessimistic, optimistic and none. To choose a locking strategy, you must consider issues such as the percentage of each type of operations you have, whether or not you use a loader and so on.

Locks are bound by transactions. You can specify the following locking settings:

- **No locking:** Running without the locking setting is the fastest. If you are using read-only data, then you might not need locking.
- **Pessimistic locking:** Acquires locks on entries, then and holds the locks until commit time. This locking strategy provides good consistency at the expense of throughput.
- **Optimistic locking:** Takes a before image of every record that the transaction touches and compares the image to the current entry values when the transaction commits. If the entry values change, then the transaction rolls back. No locks are held until commit time. This locking strategy provides better concurrency than the pessimistic strategy, at the risk of the transaction rolling back and the memory cost of making the extra copy of the entry.

Set the locking strategy on the BackingMap. You cannot change the locking strategy for each transaction. An example XML snippet that shows how to set the locking mode on a map using the XML file follows, assuming cc is the namespace for the objectgrid/config namespace:

```
<cc:backingMap name="RuntimeLifespan" lockStrategy="PESSIMISTIC" />
```

Pessimistic locking

Use the pessimistic locking strategy for read and write maps when other locking strategies are not possible. When an ObjectGrid map is configured to use the pessimistic locking strategy, a pessimistic transaction lock for a map entry is obtained when a transaction first gets the entry from the BackingMap. The

pessimistic lock is held until the application completes the transaction. Typically, the pessimistic locking strategy is used in the following situations:

- When the BackingMap is configured with or without a loader and versioning information is not available.
- When the BackingMap is used directly by an application that needs help from the eXtreme Scale for concurrency control.
- When versioning information is available, but update transactions frequently collide on the backing entries, resulting in optimistic update failures.

Because the pessimistic locking strategy has the greatest impact on performance and scalability, this strategy should only be used for read and write maps when other locking strategies are not viable. For example, these situations might include when optimistic update failures occur frequently, or when recovery from optimistic failure is difficult for an application to handle.

Optimistic locking

The optimistic locking strategy assumes that no two transactions might attempt to update the same map entry while running concurrently. Because of this belief, the lock mode does not need to be held for the life cycle of the transaction because it is unlikely that more than one transaction might update the map entry concurrently. The optimistic locking strategy is typically used in the following situations:

- When a BackingMap is configured with or without a loader and versioning information is available.
- When a BackingMap has mostly transactions that perform read operations. Insert, update, or remove operations on map entries do not occur often on the BackingMap.
- When a BackingMap is inserted, updated, or removed more frequently than it is read, but transactions rarely collide on the same map entry.

Like the pessimistic locking strategy, the methods on the ObjectMap interface determine how eXtreme Scale automatically attempts to acquire a lock mode for the map entry that is being accessed. However, the following differences between the pessimistic and optimistic strategies exist:

- Like the pessimistic locking strategy, an S lock mode is acquired by the get and getAll methods when the method is invoked. However, with optimistic locking, the S lock mode is not held until the transaction is completed. Instead, the S lock mode is released before the method returns to the application. The purpose of acquiring the lock mode is so that eXtreme Scale can ensure that only committed data from other transactions is visible to the current transaction. After eXtreme Scale has verified that the data is committed, the S lock mode is released. At commit time, an optimistic versioning check is performed to ensure that no other transaction has changed the map entry after the current transaction released its S lock mode. If an entry is not fetched from the map before it is updated, invalidated, or deleted, the eXtreme Scale run time implicitly fetches the entry from the map. This implicit get operation is performed to get the current value at the time the entry was requested to be modified.
- Unlike pessimistic locking strategy, the getForUpdate and getAllForUpdate methods are handled exactly like the get and getAll methods when the optimistic locking strategy is used. That is, an S lock mode is acquired at the start of the method and the S lock mode is released before returning to the application.

All other `ObjectMap` methods are handled exactly like they are handled for the pessimistic locking strategy. That is, when the `commit` method is invoked, an X lock mode is obtained for any map entry that is inserted, updated, removed, touched, or invalidated and the X lock mode is held until the transaction completes commit processing.

The optimistic locking strategy assumes that no concurrently running transactions attempt to update the same map entry. Because of this assumption, the lock mode does not need to be held for the life of the transaction because it is unlikely that more than one transaction might update the map entry concurrently. However, because a lock mode was not held, another concurrent transaction might potentially update the map entry after the current transaction has released its S lock mode.

To handle this possibility, eXtreme Scale gets an X lock at commit time and performs an optimistic versioning check to verify that no other transaction has changed the map entry after the current transaction read the map entry from the `BackingMap`. If another transaction changes the map entry, the version check fails and an `OptimisticCollisionException` exception occurs. This exception forces the current transaction to be rolled back and the application must try the entire transaction again. The optimistic locking strategy is very useful when a map is mostly read and it is unlikely that updates for the same map entry might occur.

No locking

When a `BackingMap` is configured to use no locking strategy, no transaction locks for a map entry are obtained.

Using no locking strategy is useful when an application is a persistence manager such as an Enterprise JavaBeans (EJB) container or when an application uses Hibernate to obtain persistent data. In this scenario, the `BackingMap` is configured without a loader and the persistence manager uses the `BackingMap` as a data cache. In this scenario, the persistence manager provides concurrency control between transactions that are accessing the same Map entries.

WebSphere eXtreme Scale does not need to obtain any transaction locks for the purpose of concurrency control. This situation assumes that the persistence manager does not release its transaction locks before updating the `ObjectGrid` map with committed changes. If the persistence manager releases its locks, then a pessimistic or optimistic lock strategy must be used. For example, suppose that the persistence manager of an EJB container is updating an `ObjectGrid` map with data that was committed in the EJB container-managed transaction. If the update of the `ObjectGrid` map occurs before the persistence manager transaction locks are released, then you can use the no lock strategy. If the `ObjectGrid` map update occurs after the persistence manager transaction locks are released, then you must use either the optimistic or pessimistic lock strategy.

Another scenario where no locking strategy can be used is when the application uses a `BackingMap` directly and a `Loader` is configured for the map. In this scenario, the loader uses the concurrency control support that is provided by a relational database management system (RDBMS) by using either Java database connectivity (JDBC) or Hibernate to access data in a relational database. The loader implementation can use either an optimistic or pessimistic approach. A loader that uses an optimistic locking or versioning approach helps to achieve the greatest amount of concurrency and performance. For more information about implementing an optimistic locking approach, see the `OptimisticCallback` section in

the information about loader considerations in the *Administration Guide*. If you are using a loader that uses pessimistic locking support of an underlying backend, you might want to use the `forUpdate` parameter that is passed on the `get` method of the Loader interface. Set this parameter to true if the `getForUpdate` method of the ObjectMap interface was used by the application to get the data. The loader can use this parameter to determine whether to request an upgradeable lock on the row that is being read. For example, DB2® obtains an upgradeable lock when an SQL select statement contains a FOR UPDATE clause. This approach offers the same deadlock prevention that is described in “Pessimistic locking” on page 131.

Distributing transactions

Use Java Message Service (JMS) for distributed transaction changes between different tiers or in environments on mixed platforms.

JMS is an ideal protocol for distributed changes between different tiers or in environments on mixed platforms. For example, some applications that use eXtreme Scale might be deployed on IBM WebSphere Application Server Community Edition, Apache Geronimo, or Apache Tomcat, whereas other applications might run on WebSphere Application Server Version 6.x. JMS is ideal for distributed changes between eXtreme Scale peers in these different environments. The high availability manager message transport is very fast, but can only distribute changes to Java virtual machines that are in a single core group. JMS is slower, but allows larger and more diverse sets of application clients to share an ObjectGrid. JMS is ideal when sharing data in an ObjectGrid between a fat Swing client and an application deployed on WebSphere Extended Deployment.

The built-in Client Invalidation Mechanism and Peer-to-Peer Replication are examples of JMS-based transactional changes distribution. See the information about configuring peer-to-peer replication with JMS in the *Administration Guide* for more information.

Implementing JMS

JMS is implemented for distributing transaction changes by using a Java object that behaves as an `ObjectGridEventListener`. This object can propagate the state in the following four ways:

1. Invalidate: Any entry that is evicted, updated or deleted is removed on all peer Java virtual machines when they receive the message.
2. Invalidate conditional: The entry is evicted only if the local version is the same or older than the version on the publisher.
3. Push: Any entry that was evicted, updated, deleted or inserted is added or overwritten on all peer Java virtual machines when they receive the JMS message.
4. Push conditional: The entry is only updated or added on the receive side if the local entry is less recent than the version that is being published.

Listen for changes for publishing

The plug-in implements the `ObjectGridEventListener` interface to intercept the `transactionEnd` event. When eXtreme Scale invokes this method, the plug-in attempts to convert the `LogSequence` list for each map that is touched by the transaction to a JMS message and then publish it. The plug-in can be configured to publish changes for all maps or a subset of maps. `LogSequence` objects are processed for the maps that have publishing enabled. The `LogSequenceTransformer` ObjectGrid class serializes a filtered `LogSequence` for

each map to a stream. After all LogSequences are serialized to the stream, then a JMS ObjectMessage is created and published to a well-known topic.

Listen for JMS messages and apply them to the local ObjectGrid

The same plug-in also starts a thread that spins in a loop, receiving all messages that are published to the well known topic. When a message arrives, it passes the message contents to the LogSequenceTransformer class where it is converted to a set of LogSequence objects. Then, a no-write-through transaction is started. Each LogSequence object is provided to the Session.processLogSequence method, which updates the local Maps with the changes. The processLogSequence method understands the distribution mode. The transaction is committed and the local cache now reflects the changes. For more information about using JMS to distribute transaction changes, see the information about distributing changes between peer Java Virtual Machines in the *Administration Guide*.

Single-partition and cross-data-grid transactions

The major distinction between WebSphere eXtreme Scale and traditional data storage solutions like relational databases or in-memory databases is the use of partitioning, which allows the cache to scale linearly. The important types of transactions to consider are single-partition and every-partition (cross-data-grid) transactions.

In general, interactions with the cache can be categorized as single-partition transactions or cross-data-grid transactions, as discussed in the following section.

Single-partition transactions

Single-partition transactions are the preferable method for interacting with caches that are hosted by WebSphere eXtreme Scale. When a transaction is limited to a single partition, then by default it is limited to a single Java virtual machine, and therefore a single server computer. A server can complete M number of these transactions per second, and if you have N computers, you can complete $M*N$ transactions per second. If your business increases and you need to perform twice as many of these transactions per second, you can double N by buying more computers. Then you can meet capacity demands without changing the application, upgrading hardware, or even taking the application offline.

In addition to letting the cache scale so significantly, single-partition transactions also maximize the availability of the cache. Each transaction only depends on one computer. Any of the other $(N-1)$ computers can fail without affecting the success or response time of the transaction. So if you are running 100 computers and one of them fails, only 1 percent of the transactions in flight at the moment that server failed are rolled back. After the server fails, WebSphere eXtreme Scale relocates the partitions that are hosted by the failed server to the other 99 computers. During this brief period, before the operation completes, the other 99 computers can still complete transactions. Only the transactions that would involve the partitions that are being relocated are blocked. After the failover process is complete, the cache can continue running, fully operational, at 99 percent of its original throughput capacity. After the failed server is replaced and returned to the data grid, the cache returns to 100 percent throughput capacity.

Cross-data-grid transactions

In terms of performance, availability and scalability, cross-data-grid transactions are the opposite of single-partition transactions. Cross-data-grid transactions access

every partition and therefore every computer in the configuration. Each computer in the data grid is asked to look up some data and then return the result. The transaction cannot complete until every computer has responded, and therefore the throughput of the entire data grid is limited by the slowest computer. Adding computers does not make the slowest computer faster and therefore does not improve the throughput of the cache.

Cross-data-grid transactions have a similar effect on availability. Extending the previous example, if you are running 100 servers and one server fails, then 100 percent of the transactions that are in progress at the moment that server failed are rolled back. After the server fails, WebSphere eXtreme Scale starts to relocate the partitions that are hosted by that server to the other 99 computers. During this time, before the failover process completes, the data grid cannot process any of these transactions. After the failover process is complete, the cache can continue running, but at reduced capacity. If each computer in the data grid serviced 10 partitions, then 10 of the remaining 99 computers receive at least one extra partition as part of the failover process. Adding an extra partition increases the workload of that computer by at least 10 percent. Because the throughput of the data grid is limited to the throughput of the slowest computer in a cross-data-grid transaction, on average, the throughput is reduced by 10 percent.

Single-partition transactions are preferable to cross-data-grid transactions for scaling out with a distributed, highly available, object cache like WebSphere eXtreme Scale. Maximizing the performance of these kinds of systems requires the use of techniques that are different from traditional relational methodologies, but you can turn cross-data-grid transactions into scalable single-partition transactions.

Best practices for building scalable data models

The best practices for building scalable applications with products like WebSphere eXtreme Scale include two categories: foundational principles and implementation tips. Foundational principles are core ideas that need to be captured in the design of the data itself. An application that does not observe these principles is unlikely to scale well, even for its mainline transactions. Implementation tips are applied for problematic transactions in an otherwise well-designed application that observes the general principles for scalable data models.

Foundational principles

Some of the important means of optimizing scalability are basic concepts or principles to keep in mind.

Duplicate instead of normalizing

The key thing to remember about products like WebSphere eXtreme Scale is that they are designed to spread data across a large number of computers. If the goal is to make most or all transactions complete on a single partition, then the data model design needs to ensure that all the data the transaction might need is located in the partition. Most of the time, the only way to achieve this is by duplicating data.

For example, consider an application like a message board. Two very important transactions for a message board are showing all the posts from a given user and all the posts on a given topic. First consider how these transactions would work with a normalized data model that contains a user record, a topic record, and a post record that contains the actual text. If posts are partitioned with user records, then displaying the topic

becomes a cross-grid transaction, and vice versa. Topics and users cannot be partitioned together because they have a many-to-many relationship.

The best way to make this message board scale is to duplicate the posts, storing one copy with the topic record and one copy with the user record. Then, displaying the posts from a user is a single-partition transaction, displaying the posts on a topic is a single-partition transaction, and updating or deleting a post is a two-partition transaction. All three of these transactions will scale linearly as the number of computers in the data grid increases.

Scalability rather than resources

The biggest obstacle to overcome when considering denormalized data models is the impact that these models have on resources. Keeping two, three, or more copies of some data can seem to use too many resources to be practical. When you are confronted with this scenario, remember the following facts: Hardware resources get cheaper every year. Second, and more importantly, WebSphere eXtreme Scale eliminates most hidden costs associated with deploying more resources.

Measure resources in terms of cost rather than computer terms such as megabytes and processors. Data stores that work with normalized relational data generally need to be located on the same computer. This required collocation means that a single larger enterprise computer needs to be purchased rather than several smaller computers. With enterprise hardware, it is not uncommon for one computer to be capable of completing one million transactions per second to cost much more than the combined cost of 10 computers capable of doing 100,000 transactions per second each.

A business cost in adding resources also exists. A growing business eventually runs out of capacity. When you run out of capacity, you either need to shut down while moving to a bigger, faster computer, or create a second production environment to which you can switch. Either way, additional costs will come in the form of lost business or maintaining almost twice the capacity needed during the transition period.

With WebSphere eXtreme Scale, the application does not need to be shut down to add capacity. If your business projects that you need 10 percent more capacity for the coming year, then increase the number of computers in the data grid by 10 percent. You can increase this percentage without application downtime and without purchasing excess capacity.

Avoid data transformations

When you are using WebSphere eXtreme Scale, data should be stored in a format that is directly consumable by the business logic. Breaking the data down into a more primitive form is costly. The transformation needs to be done when the data is written and when the data is read. With relational databases this transformation is done out of necessity, because the data is ultimately persisted to disk quite frequently, but with WebSphere eXtreme Scale, you do not need to perform these transformations. For the most part data is stored in memory and can therefore be stored in the exact form that the application needs.

Observing this simple rule helps denormalize your data in accordance with the first principle. The most common type of transformation for business data is the JOIN operations that are necessary to turn normalized data into

a result set that fits the needs of the application. Storing the data in the correct format implicitly avoids performing these JOIN operations and produces a denormalized data model.

Eliminate unbounded queries

No matter how well you structure your data, unbounded queries do not scale well. For example, do not have a transaction that asks for a list of all items sorted by value. This transaction might work at first when the total number of items is 1000, but when the total number of items reaches 10 million, the transaction returns all 10 million items. If you run this transaction, the two most likely outcomes are the transaction timing out, or the client encountering an out-of-memory error.

The best option is to alter the business logic so that only the top 10 or 20 items can be returned. This logic alteration keeps the size of the transaction manageable no matter how many items are in the cache.

Define schema

The main advantage of normalizing data is that the database system can take care of data consistency behind the scenes. When data is denormalized for scalability, this automatic data consistency management no longer exists. You must implement a data model that can work in the application layer or as a plug-in to the distributed data grid to guarantee data consistency.

Consider the message board example. If a transaction removes a post from a topic, then the duplicate post on the user record needs to be removed. Without a data model, it is possible a developer would write the application code to remove the post from the topic and forget to remove the post from the user record. However, if the developer were using a data model instead of interacting with the cache directly, the `removePost` method on the data model could pull the user ID from the post, look up the user record, and remove the duplicate post behind the scenes.

Alternately, you can implement a listener that runs on the actual partition that detects the change to the topic and automatically adjusts the user record. A listener might be beneficial because the adjustment to the user record could happen locally if the partition happens to have the user record, or even if the user record is on a different partition, the transaction takes place between servers instead of between the client and server. The network connection between servers is likely to be faster than the network connection between the client and the server.

Avoid contention

Avoid scenarios such as having a global counter. The data grid will not scale if a single record is being used a disproportionate number of times compared to the rest of the records. The performance of the data grid will be limited by the performance of the computer that holds the given record.

In these situations, try to break the record up so it is managed per partition. For example consider a transaction that returns the total number of entries in the distributed cache. Instead of having every insert and remove operation access a single record that increments, have a listener on each partition track the insert and remove operations. With this listener tracking, insert and remove can become single-partition operations.

Reading the counter will become a cross-data-grid operation, but for the most part, it was already as inefficient as a cross-data-grid operation because its performance was tied to the performance of the computer hosting the record.

Implementation tips

You can also consider the following tips to achieve the best scalability.

Use reverse-lookup indexes

Consider a properly denormalized data model where customer records are partitioned based on the customer ID number. This partitioning method is the logical choice because nearly every business operation performed with the customer record uses the customer ID number. However, an important transaction that does not use the customer ID number is the login transaction. It is more common to have user names or e-mail addresses for login instead of customer ID numbers.

The simple approach to the login scenario is to use a cross-data-grid transaction to find the customer record. As explained previously, this approach does not scale.

The next option might be to partition on user name or e-mail. This option is not practical because all the customer ID based operations become cross-data-grid transactions. Also, the customers on your site might want to change their user name or e-mail address. Products like WebSphere eXtreme Scale need the value that is used to partition the data to remain constant.

The correct solution is to use a reverse lookup index. With WebSphere eXtreme Scale, a cache can be created in the same distributed grid as the cache that holds all the user records. This cache is highly available, partitioned and scalable. This cache can be used to map a user name or e-mail address to a customer ID. This cache turns login into a two partition operation instead of a cross-grid operation. This scenario is not as good as a single-partition transaction, but the throughput still scales linearly as the number of computers increases.

Compute at write time

Commonly calculated values like averages or totals can be expensive to produce because these operations usually require reading a large number of entries. Because reads are more common than writes in most applications, it is efficient to compute these values at write time and then store the result in the cache. This practice makes read operations both faster and more scalable.

Optional fields

Consider a user record that holds a business, home, and telephone number. A user could have all, none or any combination of these numbers defined. If the data were normalized then a user table and a telephone number table would exist. The telephone numbers for a given user could be found using a JOIN operation between the two tables.

De-normalizing this record does not require data duplication, because most users do not share telephone numbers. Instead, empty slots in the user record must be allowed. Instead of having a telephone number table, add three attributes to each user record, one for each telephone number type.

This addition of attributes eliminates the JOIN operation and makes a telephone number lookup for a user a single-partition operation.

Placement of many-to-many relationships

Consider an application that tracks products and the stores in which the products are sold. A single product is sold in many stores, and a single store sells many products. Assume that this application tracks 50 large retailers. Each product is sold in a maximum of 50 stores, with each store selling thousands of products.

Keep a list of stores inside the product entity (arrangement A), instead of keeping a list of products inside each store entity (arrangement B). Looking at some of the transactions this application would have to perform illustrates why arrangement A is more scalable.

First look at updates. With arrangement A, removing a product from the inventory of a store locks the product entity. If the data grid holds 10000 products, only 1/10000 of the grid needs to be locked to perform the update. With arrangement B, the data grid only contains 50 stores, so 1/50 of the grid must be locked to complete the update. So even though both of these could be considered single-partition operations, arrangement A scales out more efficiently.

Now, considering reads with arrangement A, looking up the stores at which a product is sold is a single-partition transaction that scales and is fast because the transaction only transmits a small amount of data. With arrangement B, this transaction becomes a cross-data-grid transaction because each store entity must be accessed to see if the product is sold at that store, which reveals an enormous performance advantage for arrangement A.

Scaling with normalized data

One legitimate use of cross-data-grid transactions is to scale data processing. If a data grid has 5 computers and a cross-data-grid transaction is dispatched that sorts through about 100,000 records on each computer, then that transaction sorts through 500,000 records. If the slowest computer in the data grid can perform 10 of these transactions per second, then the data grid is capable of sorting through 5,000,000 records per second. If the data in the grid doubles, then each computer must sort through 200,000 records, and each transaction sorts through 1,000,000 records. This data increase decreases the throughput of the slowest computer to 5 transactions per second, thereby reducing the throughput of the data grid to 5 transactions per second. Still, the data grid sorts through 5,000,000 records per second.

In this scenario, doubling the number of computer allows each computer to return to its previous load of sorting through 100,000 records, allowing the slowest computer to process 10 of these transactions per second. The throughput of the data grid stays the same at 10 requests per second, but now each transaction processes 1,000,000 records, so the grid has doubled its capacity in terms of processing records to 10,000,000 per second.

Applications such as a search engine that need to scale both in terms of data processing to accommodate the increasing size of the Internet and throughput to accommodate growth in the number of users, you must create multiple data grids, with a round robin of the requests between the grids. If you need to scale up the throughput, add computers and add

another data grid to service requests. If data processing needs to be scaled up, add more computers and keep the number of data grids constant.

Using locking

Locks have life cycles and different types of locks are compatible with others in various ways. Locks must be handled in the correct order to avoid deadlock scenarios.

Locks

Locks have life cycles and different types of locks are compatible with others in various ways. Locks must be handled in the correct order to avoid deadlock scenarios.

Shared, upgradeable, and exclusive locks

When an application calls any method of the ObjectMap interface, uses the find methods on an index, or does a query, eXtreme Scale automatically attempts to acquire a lock for the map entry that is being accessed. WebSphere eXtreme Scale uses the following lock modes based on the method the application calls in the ObjectMap interface.

- The get and getAll methods on the ObjectMap interface, index methods, and queries acquire an *S lock*, or a shared lock mode for the key of a map entry. The duration that the S lock is held depends on the transaction isolation level used. An S lock mode allows concurrency between transactions that attempt to acquire an S or an upgradeable lock (U lock) mode for the same key, but blocks other transactions that attempt to get an exclusive lock (X lock) mode for the same key.
- The getForUpdate and getAllForUpdate methods acquire a *U lock*, or an upgradeable lock mode for the key of a map entry. The U lock is held until the transaction completes. A U lock mode allows concurrency between transactions that acquire an S lock mode for the same key, but blocks other transactions that attempt to acquire a U lock or X lock mode for the same key.
- The put, putAll, remove, removeAll, insert, update, and touch acquire an *X lock*, or exclusive lock mode for the key of a map entry. The X lock is held until the transaction completes. An X lock mode ensures that only one transaction is inserting, updating, or removing a map entry of a given key value. An X lock blocks all other transactions that attempt to acquire a S, U, or X lock mode for the same key.
- The global invalidate and global invalidateAll methods acquire an X lock for each map entry that is invalidated. The X lock is held until the transaction completes. No locks are acquired for the local invalidate and local invalidateAll methods because none of the BackingMap entries are invalidated by local invalidate method calls.

From the preceding definitions, it is obvious that an S lock mode is weaker than a U lock mode because it allows more transactions to run concurrently when accessing the same map entry. The U lock mode is slightly stronger than the S lock mode because it blocks other transactions that are requesting either a U or X lock mode. The S lock mode only blocks other transactions that are requesting an X lock mode. This small difference is important in preventing some deadlocks from occurring. The X lock mode is the strongest lock mode because it blocks all other transactions attempting to get an S, U, or X lock mode for the same map entry. The net effect of an X lock mode is to ensure that only one transaction can insert, update, or remove a map entry and to prevent updates from being lost when more than one transaction is attempting to update the same map entry.

The following table is a lock mode compatibility matrix that summarizes the described lock modes, which you can use to determine which lock modes are compatible with each other. To read this matrix, the row in the matrix indicates a lock mode that is already granted. The column indicates the lock mode that is requested by another transaction. If Yes is displayed in the column, then the lock mode requested by the other transaction is granted because it is compatible with the lock mode that is already granted. No indicates that the lock mode is not compatible and the other transaction must wait for the first transaction to release the lock that it owns.

Table 4. Lock mode compatibility matrix

| Lock | Lock type S (shared) | Lock type U (upgradeable) | Lock type X (exclusive) | Strength |
|-----------------|----------------------|---------------------------|-------------------------|-----------|
| S (shared) | Yes | Yes | No | weakest |
| U (upgradeable) | Yes | No | No | normal |
| X (exclusive) | No | No | No | strongest |

Locking deadlocks

Consider the following sequence of lock mode requests:

1. X lock is granted to transaction 1 for key1.
2. X lock is granted to transaction 2 for key2.
3. X lock requested by transaction 1 for key2. (Transaction 1 blocks waiting for lock owned by transaction 2.)
4. X lock requested by transaction 2 for key1. (Transaction 2 blocks waiting for lock owned by transaction 1.)

The preceding sequence is the classic deadlock example of two transactions that attempt to acquire more than a single lock, and each transaction acquires the locks in a different order. To prevent this deadlock, each transaction must obtain the multiple locks in the same order. If the OPTIMISTIC lock strategy is used and the flush method on the ObjectMap interface is never used by the application, then lock modes are requested by the transaction only during the commit cycle. During the commit cycle, eXtreme Scale determines the keys for the map entries that need to be locked and requests the lock modes in key sequence (deterministic behavior). With this method, eXtreme Scale prevents the large majority of the classic deadlocks. However, eXtreme Scale does not and cannot prevent all possible deadlock scenarios. A few scenarios exist that the application needs to consider. Following are the scenarios that the application must be aware of and take preventative action against.

One scenario exists where eXtreme Scale is able to detect a deadlock without having to wait for a lock wait timeout to occur. If this scenario does occur, a `com.ibm.websphere.objectgrid.LockDeadlockException` exception results. Consider the following code snippet:

```
Session sess = ...;
ObjectMap person = sess.getMap("PERSON");
sess.begin();
Person p = (IPerson)person.get("Lynn");
// Lynn had a birthday, so we make her 1 year older.
p.setAge( p.getAge() + 1 );
person.put( "Lynn", p );
sess.commit();
```

In this situation, Lynn's boyfriend wants to make her older than she is now, and both Lynn and her boyfriend run this transaction concurrently. In this situation,

both transactions own an S lock mode on the Lynn entry of the PERSON map as a result of the `person.get("Lynn")` method invocation. As a result of the `person.put("Lynn", p)` method call, both transactions attempt to upgrade the S lock mode to an X lock mode. Both transactions block waiting for the other transaction to release the S lock mode it owns. As a result, a deadlock occurs because a circular wait condition exists between the two transactions. A circular wait condition results when more than one transaction attempts to promote a lock from a weaker to a stronger mode for the same map entry. In this scenario, a `LockDeadlockException` exception results instead of a `LockTimeoutException` exception.

The application can prevent the `LockDeadlockException` exception for the preceding example by using the optimistic lock strategy instead of the pessimistic lock strategy. Using the optimistic lock strategy is the preferred solution when the map is mostly read and updates to the map are infrequent. If the pessimistic lock strategy must be used, the `getForUpdate` method can be used instead of the `get` method in the above example or a transaction isolation level of `TRANSACTION_READ_COMMITTED` can be used.

For more information, see the topic on locking strategies in the *Product Overview*.

Using the `TRANSACTION_READ_COMMITTED` transaction isolation level prevents the S lock that is acquired by the `get` method from being held until the transaction completes. If the key is never invalidated in the transactional cache, repeatable reads are still guaranteed.

See the topic on map entry locking in the *Administration Guide* for more information.

An alternative to changing the transaction isolation level is to use the `getForUpdate` method. The first transaction to call the `getForUpdate` method acquires a U lock mode instead of an S lock. This lock mode causes the second transaction to block when it calls the `getForUpdate` method because only one transaction is granted a U lock mode. Because the second transaction is blocked, it does not own any lock mode on the Lynn map entry. The first transaction does not block when it attempts to upgrade the U lock mode to an X lock mode as a result of the `put` method call from the first transaction. This feature demonstrates why U lock mode is called the *upgradeable* lock mode. When the first transaction is completed, the second transaction unblocks and is granted the U lock mode. An application can prevent the lock promotion deadlock scenario by using the `getForUpdate` method instead of the `get` method when pessimistic lock strategy is being used.

Important: This solution does not prevent read-only transactions from being able to read a map entry. Read-only transactions call the `get` method, but never call the `put`, `insert`, `update`, or `remove` methods. Concurrency is just as high as when the regular `get` method is used. The only reduction in concurrency occurs when the `getForUpdate` method is called by more than one transaction for the same map entry.

You must be aware when a transaction calls the `getForUpdate` method on more than one map entry to ensure that the U locks are acquired in the same order by each transaction. For example, suppose that the first transaction calls the `getForUpdate` method for the key 1 and the `getForUpdate` method for key 2. Another concurrent transaction calls the `getForUpdate` method for the same keys, but in reverse order. This sequence causes the classic deadlock because multiple locks are obtained in different orders by different transactions. The application still

needs to ensure that every transaction accesses multiple map entries in key sequence to ensure that deadlock does not occur. Because the U lock is obtained at the time that the `getForUpdate` method is called rather than at commit time, the eXtreme Scale cannot order the lock requests like it does during the commit cycle. The application must control the lock ordering in this case.

Using the `flush` method on the `ObjectMap` interface before a commit can introduce additional lock ordering considerations. The `flush` method is typically used to force changes made to the map out to the backend through the `Loader` plug-in. In this situation, the backend uses its own lock manager to control concurrency, so the lock wait condition and deadlock can occur in backend rather than in the eXtreme Scale lock manager. Consider the following transaction:

```
Session sess = ...;
ObjectMap person = sess.getMap("PERSON");
boolean activeTran = false;
try
{
    sess.begin();
    activeTran = true;
    Person p = (IPerson)person.get("Lynn");
    p.setAge( p.getAge() + 1 );
    person.put( "Lynn", p );
    person.flush();
    ...
    p = (IPerson)person.get("Tom");
    p.setAge( p.getAge() + 1 );
    sess.commit();
    activeTran = false;
}
finally
{
    if ( activeTran ) sess.rollback();
}
```

Suppose that another transaction also updated the Tom person, called the `flush` method, and then updated the Lynn person. If this situation occurred, the following interleaving of the two transactions results in a database deadlock condition:

```
X lock is granted to transaction 1 for "Lynn" when flush is executed.
X lock is granted to transaction 2 for "Tom" when flush is executed..
X lock requested by transaction 1 for "Tom" during commit processing.
(Transaction 1 blocks waiting for lock owned by transaction 2.)
X lock requested by transaction 2 for "Lynn" during commit processing.
(Transaction 2 blocks waiting for lock owned by transaction 1.)
```

This example demonstrates that the use of the `flush` method can cause a deadlock to occur in the database rather than in eXtreme Scale. This deadlock example can occur regardless of what lock strategy is used. The application must take care to prevent this kind of deadlock from occurring when using the `flush` method and when a `Loader` is plugged into the `BackingMap`. The preceding example also illustrates another reason why eXtreme Scale has a lock wait timeout mechanism. A transaction that is waiting for a database lock might be waiting while it owns an eXtreme Scale map entry lock. Consequently, problems at database level can cause excessive wait times for an eXtreme Scale lock mode and result in a `LockTimeoutException` exception.

Implementing exception handling in locking scenarios

To prevent locks from being held for excessive amounts of time when a `LockTimeoutException` exception or a `LockDeadlockException` exception occurs, an

application must ensure that it catches unexpected exceptions and calls the rollback method when something unexpected occurs.

Procedure

1. Catch the exception, and display resulting message.

```
try {  
    ...  
} catch (ObjectGridException oe) {  
    System.out.println(oe);  
}
```

The following exception displays as a result:

```
com.ibm.websphere.objectgrid.plugins.LockDeadlockException: Message
```

This message represents the string that is passed as a parameter when the exception is created and thrown.

2. Roll back the transaction after an exception:

```
Session sess = ...;  
ObjectMap person = sess.getMap("PERSON");  
boolean activeTran = false;  
try  
{  
    sess.begin();  
    activeTran = true;  
    Person p = (IPerson)person.get("Lynn");  
    // Lynn had a birthday, so we make her 1 year older.  
    p.setAge( p.getAge() + 1 );  
    person.put( "Lynn", p );  
    sess.commit();  
    activeTran = false;  
}  
finally  
{  
    if ( activeTran ) sess.rollback();  
}
```

The finally block in the snippet of code ensures that a transaction is rolled back when an unexpected exception occurs. It not only handles a LockDeadlockException exception, but any other unexpected exception that might occur. The finally block handles the case where an exception occurs during a commit method invocation. This example is not the only way to deal with unexpected exceptions, and there might be cases where an application wants to catch some of the unexpected exceptions that can occur and display one of its application exceptions. You can add catch blocks as appropriate, but the application must ensure that the snippet of code does not exit without completing the transaction.

Configuring a locking strategy

You can define an optimistic, a pessimistic, or no locking strategy on each BackingMap in the WebSphere eXtreme Scale configuration.

About this task

Each BackingMap instance can be configured to use one of the following locking strategies:

1. Optimistic locking mode
2. Pessimistic locking mode
3. None

The default lock strategy is OPTIMISTIC. Use optimistic locking when data is changed infrequently. Locks are only held for a short duration while data is being read from the cache and copied to the transaction. When the transaction cache is synchronized with the main cache, any cache objects that have been updated are checked against the original version. If the check fails, then the transaction is rolled back and an `OptimisticCollisionException` exception results.

The PESSIMISTIC lock strategy acquires locks for cache entries and should be used when data is changed frequently. Any time a cache entry is read, a lock is acquired and conditionally held until the transaction completes. The duration of some locks can be tuned using transaction isolation levels for the session.

If locking is not required because the data is never updated or is only updated during quiet periods, you can disable locking by using the NONE lock strategy. This strategy is very fast because a lock manager is not required. The NONE lock strategy is ideal for look-up tables or read-only maps.

For more information about locking strategies, see the information about locking strategies in the *Product Overview*.

You can specify a locking strategy programmatically or with XML. For more information about locking, see the information about locking strategies in the *Product Overview*.

Procedure

• Configure an optimistic locking strategy

- Programmatically using the `setLockStrategy` method:

```
import com.ibm.websphere.objectgrid.BackingMap;
import com.ibm.websphere.objectgrid.LockStrategy;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
...
ObjectGrid og =
    ObjectGridManagerFactory.getObjectGridManager().createObjectGrid("test");
BackingMap bm = og.defineMap("optimisticMap");
bm.setLockStrategy( LockStrategy.OPTIMISTIC );
```

- Using the `lockStrategy` attribute in the ObjectGrid descriptor XML file:

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
    xmlns="http://ibm.com/ws/objectgrid/config">
  <objectGrids>
    <objectGrid name="test">
      <backingMap name="optimisticMap"
        lockStrategy="OPTIMISTIC"/>
    </objectGrid>
  </objectGrids>
</objectGridConfig>
```

• Configure a pessimistic locking strategy

- Programmatically using the `setLockStrategy` method:

```
specify pessimistic strategy programmatically
import com.ibm.websphere.objectgrid.BackingMap;
import com.ibm.websphere.objectgrid.LockStrategy;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
...
ObjectGrid og =
```

```

    ObjectGridManagerFactory.getObjectGridManager().createObjectGrid("test");
    BackingMap bm = og.defineMap("pessimisticMap");
    bm.setLockStrategy( LockStrategy.PESSIMISTIC);

```

- Using the lockStrategy attribute in the ObjectGrid descriptor XML file:

specify pessimistic strategy using XML

```

<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">

  <objectGrids>
    <objectGrid name="test">
      <backingMap name="pessimisticMap"
        lockStrategy="PESSIMISTIC"/>
    </objectGrid>
  </objectGrids>
</objectGridConfig>

```

- **Configure a no locking strategy**

- Programmatically using the setLockStrategy method:

```

import com.ibm.websphere.objectgrid.BackingMap;
import com.ibm.websphere.objectgrid.LockStrategy;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
...
ObjectGrid og =
  ObjectGridManagerFactory.getObjectGridManager().createObjectGrid("test");
BackingMap bm = og.defineMap("noLockingMap");
bm.setLockStrategy( LockStrategy.NONE);

```

- Using the lockStrategy attribute in the ObjectGrid descriptor XML file:

```

<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">

  <objectGrids>
    <objectGrid name="test">
      <backingMap name="noLockingMap"
        lockStrategy="NONE"/>
    </objectGrid>
  </objectGrids>
</objectGridConfig>

```

What to do next

To avoid a java.lang.IllegalStateException exception, you must call the setLockStrategy method before calling the initialize or getSession methods on the ObjectGrid instance.

Configuring the lock timeout value

The lock timeout value on a BackingMap instance is used to ensure that an application does not wait endlessly for a lock mode to be granted because of a deadlock condition that occurs due to an application error.

Before you begin

To configure the lock timeout value, the locking strategy must be set to either OPTIMISTIC or PESSIMISTIC. See “Configuring a locking strategy” on page 145 for more information.

About this task

When a `LockTimeoutException` exception occurs, the application must determine if the timeout is occurring because the application is running slower than expected, or if the timeout occurred because of a deadlock condition. If an actual deadlock condition occurred, then increasing the lock wait timeout value does not eliminate the exception. Increasing the timeout results in the exception taking longer to occur. However, if increasing the lock wait timeout value does eliminate the exception, then the problem occurred because the application was running slower than expected. The application in this case must determine why performance is slow.

To prevent deadlocks from occurring, the lock manager has a default timeout value of 15 seconds. If the timeout limit is exceeded, a `LockTimeoutException` exception occurs. If your system is heavily loaded, the default timeout value might cause the `LockTimeoutException` exceptions to occur when no deadlock exists. In this situation, you can increase the lock timeout value programmatically or in the ObjectGrid descriptor XML file.

Procedure

- Configure a lock timeout value programmatically on a `BackingMap` instance with the `setLockTimeout` method.

The following example illustrates how to set the lock wait timeout value for the `map1` backing map to 60 seconds:

```
import com.ibm.websphere.objectgrid.BackingMap;
import com.ibm.websphere.objectgrid.LockStrategy;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
...
ObjectGrid og =
    ObjectGridManagerFactory.getObjectGridManager().createObjectGrid("test");
BackingMap bm = og.defineMap("map1");
bm.setLockStrategy( LockStrategy.PESSIMISTIC );
bm.setLockTimeout( 60 );
```

To avoid a `java.lang.IllegalStateException` exception, call both the `setLockStrategy` method and the `setLockTimeout` method before calling either the `initialize` or `getSession` methods on the `ObjectGrid` instance. The `setLockTimeout` method parameter is a Java primitive integer that specifies the number of seconds that eXtreme Scale waits for a lock mode to be granted. If a transaction waits longer than the lock wait timeout value configured for the `BackingMap`, a `com.ibm.websphere.objectgrid.LockTimeoutException` exception results.

- Configure the lock timeout value using the `lockTimeout` attribute in the ObjectGrid descriptor XML file.

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
    xmlns="http://ibm.com/ws/objectgrid/config">
    <objectGrids>
        <objectGrid name="test">
            <backingMap name="optimisticMap"
                lockStrategy="OPTIMISTIC"
                lockTimeout="60"/>
```

- Override the lock wait timeout for a single `ObjectMap` instance. Use the `ObjectMap.setLockTimeout` method to override the lock timeout value for a

specific ObjectMap instance. The lock timeout value affects all transactions started after the new timeout value is set. This method can be useful when lock collisions are possible or expected in select transactions.

Locking performance best practices

Locking strategies and transaction isolation settings affect the performance of your applications.

Retrieve a cached instance

Pessimistic locking strategy

Use the pessimistic locking strategy for read and write map operations where keys often collide. The pessimistic locking strategy has the greatest impact on performance.

Read committed and read uncommitted transaction isolation

When you are using pessimistic locking strategy, set the transaction isolation level using the `Session.setTransactionIsolation` method. For read committed or read uncommitted isolation, use the `Session.TRANSACTION_READ_COMMITTED` or `Session.TRANSACTION_READ_UNCOMMITTED` arguments depending on the isolation. To reset the transaction isolation level to the default pessimistic locking behavior, use the `Session.setTransactionIsolation` method with the `Session.REPEATABLE_READ` argument.

Read committed isolation reduces the duration of shared locks, which can improve concurrency and reduce the chance for deadlocks. This isolation level should be used when a transaction does not need assurances that read values remain unchanged for the duration of the transaction.

Use an uncommitted read when the transaction does not need to see the committed data.

Optimistic locking strategy

Optimistic locking is the default configuration. This strategy improves both performance and scalability compared to the pessimistic strategy. Use this strategy when your applications can tolerate some optimistic update failures, while still performing better than the pessimistic strategy. This strategy is excellent for read operations and infrequent update applications.

OptimisticCallback plug-in

The optimistic locking strategy makes a copy of the cache entries and compares them as needed. This operation can be expensive because copying the entry might involve cloning or serialization. To implement the fastest possible performance, implement the custom plug-in for non-entity maps.

See for more information. See the information about the OptimisticCallback plug-in in the *Product Overview* for more information.

Use version fields for entities

When you are using optimistic locking with entities, use the `@Version` annotation or the equivalent attribute in the Entity metadata descriptor file. The version

annotation gives the ObjectGrid a very efficient way of tracking the version of an object. If the entity does not have a version field and optimistic locking is used for the entity, then the entire entity must be copied and compared.

None locking strategy

Use the none locking strategy for applications that are read only. The none locking strategy does not obtain any locks or use a lock manager. Therefore, this strategy offers the most concurrency, performance and scalability.

Map entry locks with query and indexes

This topic describes how eXtreme Scale Query APIs and the MapRangeIndex indexing plug-in interact with locks and some best practices to increase concurrency and decrease deadlocks when using the pessimistic locking strategy for maps.

Overview

The ObjectGrid Query API allows SELECT queries over ObjectMap cache objects and entities. When a query is run, the query engine uses a MapRangeIndex when possible to find matching keys that match values in the query's WHERE clause or to bridge relationships. When an index isn't available, the query engine will scan each entry in one or more maps to find the appropriate entries. Both the query engine and index plugins will acquire locks to verify consistent data, depending on the locking strategy, transaction isolation level, and transaction state.

Locking with the HashIndex plug-in

The eXtreme Scale HashIndex plug-in allows finding keys based on a single attribute stored in the cache entry value. The index stores the indexed value in a separate data structure from the cache map. The index validates the keys against map entries before returning to the user to try to achieve an accurate result set. When the pessimistic lock strategy is used and the index is used against a local ObjectMap instance (versus a client/server ObjectMap), the index will acquire locks for each matching entry. When using optimistic locking or a remote ObjectMap, the locks are always immediately released.

The type of lock that is acquired depends upon the forUpdate argument passed to the ObjectMap.getIndex method. The forUpdate argument specifies the type of lock that the index should acquire. If false, a shareable (S) lock is acquired and if true, an upgradeable (U) lock is acquired.

If the lock type is shareable, the transaction isolation setting for the session is applied and affects the duration of the lock. See the transaction isolation topic for details on how transaction isolation is used to add concurrency to applications.

Shared locks with query

The eXtreme Scale query engine acquires S locks when needed to introspect the cache entries to discover if they satisfy the query's filter criteria. When using repeatable read transaction isolation with pessimistic locking, the S locks are only retained for the elements that are included in the query result and are released for any entries that are not included in the result. If using a lower transaction isolation level or optimistic locking, the S locks are not retained.

Shared locks with client to server query

When using the eXtreme Scale query from a client, the query typically runs on the server unless all of the maps or entities referenced in the query are local to the client (for example: a client-replicated map or a query result entity). All queries that run in a read/write transaction will retain S locks as described in the previous section. If the transaction is not a read/write transaction, then a session is not retained on the server and the S locks are released.

A read/write transaction is only routed to a primary partition and a session is maintained on the server for the client session. A transaction can be promoted to read/write under the following conditions:

1. Any map configured to use pessimistic locking is accessed using the `ObjectMap.get` and `getAll` API methods or the `EntityManager.find` methods.
2. The transaction is flushed, causing updates to be sent to the server.
3. Any map configured to use optimistic locking is accessed using the `ObjectMap.getForUpdate` or `EntityManager.findForUpdate` method.

Upgradeable locks with query

Shareable locks are useful when concurrency and consistency is important. It guarantees that an entry's value does not change for the life of the transaction. No other transaction can change the value while any other S locks are held, and only one other transaction can establish an intent to update the entry. See the [Pessimistic Locking Mode](#) topic for details on the S, U and X locking modes.

Upgradeable locks are used to identify the intent to update a cache entry when using the pessimistic lock strategy. It allows synchronization between transactions that want to modify a cache entry. Transactions can still view the entry using an S lock, but other transactions are prevented from acquiring a U lock or an X lock. In many scenarios, acquiring a U lock without first acquiring an S lock is necessary to avoid deadlocks. See the [Pessimistic Locking Mode](#) topic for common deadlock examples.

The `ObjectQuery` and `EntityManager Query` interfaces provide the `setForUpdate` method to identify the intended use for the query result. Specifically, the query engine acquires U locks instead of S locks for each map entry involved in the query result:

```
ObjectMap orderMap = session.getMap("Order");
ObjectQuery q = session.createQuery("SELECT o FROM Order o WHERE o.orderDate=?1");
q.setParameter(1, "20080101");
q.setForUpdate(true);
session.begin();
// Run the query. Each order has U lock
Iterator result = q.getResultIterator();
// For each order, update the status.
while(result.hasNext()) {
    Order o = (Order) result.next();
    o.status = "shipped";
    orderMap.update(o.getId(), o);
}
// When committed, the
session.commit();

Query q = em.createQuery("SELECT o FROM Order o WHERE o.orderDate=?1");
q.setParameter(1, "20080101");
q.setForUpdate(true);
emTran.begin();
// Run the query. Each order has U lock
```

```

Iterator result = q.getResultIterator();
// For each order, update the status.
while(result.hasNext()) {
    Order o = (Order) result.next();
    o.status = "shipped";
}
tmTran.commit();

```

When the **setForUpdate** attribute is enabled, the transaction is automatically converted to a read/write transaction and the locks are held on the server as expected. If the query cannot use any indexes, then the map must be scanned which will result in temporary U locks for map entries that do not satisfy the query result, and hold U locks for entries that are included in the result.

Transaction isolation

For transactions, you can configure each backing map configuration with one of three lock strategies: pessimistic, optimistic or none. When you are using pessimistic and optimistic locking, eXtreme Scale uses shared (S), upgradeable (U) and exclusive (X) locks to maintain consistency. This locking behavior is most notable when using pessimistic locking, because optimistic locks are not held. You can use one of three transaction isolation levels to tune the locking semantics that eXtreme Scale uses to maintain consistency in each cache map: repeatable read, read committed and read uncommitted.

Transaction isolation overview

Transaction isolation defines how the changes that are made by one operation become visible to other concurrent operations.

WebSphere eXtreme Scale supports three transaction isolation levels with which you can further tune the locking semantics that eXtreme Scale uses to maintain consistency in each cache map: repeatable read, read committed and read uncommitted. The transaction isolation level is set on the Session interface using the `setTransactionIsolation` method. The transaction isolation can be changed any time during the life of the session, if a transaction is not currently in progress.

The product enforces the various transaction isolation semantics by adjusting the way in which shared (S) locks are requested and held. Transaction isolation has no effect on maps configured to use the optimistic or none locking strategies or when upgradeable (U) locks are acquired.

Repeatable read with pessimistic locking

The repeatable read transaction isolation level is the default. This isolation level prevents dirty reads and non-repeatable reads, but does not prevent phantom reads. A dirty read is a read operation that occurs on data that has been modified by a transaction but has not been committed. A non-repeatable read might occur when read locks are not acquired when performing a read operation. A phantom read can occur when two identical read operations are performed, but two different sets of results are returned because an update has occurred on the data between the read operations. The product achieve a repeatable read by holding onto any S locks until the transaction that owns the lock completes. Because an X lock is not granted until all S locks are released, all transactions holding the S lock are guaranteed to see the same value when re-read.

```

map = session.getMap("Order");
session.setTransactionIsolation(Session.TRANSACTION_REPEATABLE_READ);
session.begin();

```

```

// An S lock is requested and held and the value is copied into
// the transactional cache.
Order order = (Order) map.get("100");
// The entry is evicted from the transactional cache.
map.invalidate("100", false);

// The same value is requested again. It already holds the
// lock, so the same value is retrieved and copied into the
// transactional cache.
Order order2 (Order) = map.get("100");

// All locks are released after the transaction is synchronized
// with cache map.
session.commit();

```

Phantom reads are possible when you are using queries or indexes because locks are not acquired for ranges of data, only for the cache entries that match the index or query criteria. For example:

```

session1.setTransactionIsolation(Session.TRANSACTION_REPEATABLE_READ);
session1.begin();

// A query is run which selects a range of values.
ObjectQuery query = session1.createObjectQuery
    ("SELECT o FROM Order o WHERE o.itemName='Widget'");

// In this case, only one order matches the query filter.
// The order has a key of "100".
// The query engine automatically acquires an S lock for Order "100".
Iterator result = query.getResultIterator();

// A second transaction inserts an order that also matches the query.
Map orderMap = session2.getMap("Order");
orderMap.insert("101", new Order("101", "Widget"));

// When the query runs again in the current transaction, the
// new order is visible and will return both Orders "100" and "101".
result = query.getResultIterator();

// All locks are released after the transaction is synchronized
// with cache map.
session.commit();

```

Read committed with pessimistic locking

The read committed transaction isolation level can be used with eXtreme Scale, which prevents dirty reads, but does not prevent non-repeatable reads or phantom reads, so eXtreme Scale continues to use S locks to read data from the cache map, but immediately releases the locks.

```

map1 = session1.getMap("Order");
session1.setTransactionIsolation(Session.TRANSACTION_READ_COMMITTED);
session1.begin();

// An S lock is requested but immediately released and
//the value is copied into the transactional cache.

Order order = (Order) map1.get("100");

// The entry is evicted from the transactional cache.
map1.invalidate("100", false);

// A second transaction updates the same order.
// It acquires a U lock, updates the value, and commits.
// The ObjectGrid successfully acquires the X lock during

```

```

// commit since the first transaction is using read
// committed isolation.

Map orderMap2 = session2.getMap("Order");
session2.begin();
order2 = (Order) orderMap2.getForUpdate("100");
order2.quantity=2;
orderMap2.update("100", order2);
session2.commit();

// The same value is requested again. This time, they
// want to update the value, but it now reflects
// the new value
Order order1Copy (Order) = map1.getForUpdate("100");

```

Read uncommitted with pessimistic locking

The read uncommitted transaction isolation level can be used with eXtreme Scale, which is a level that allows dirty reads, non-repeatable reads and phantom reads.

Optimistic collision exception

You can receive an `OptimisticCollisionException` directly, or receive it with an `ObjectGridException`.

The following code is an example of how to catch the exception and then display its message:

```

try {
...
} catch (ObjectGridException oe) {
    System.out.println(oe);
}

```

Exception cause

`OptimisticCollisionException` is created in a situation in which two different clients try to update the same map entry at relatively the same time. For example, if one client attempts to commit a session and update the map entry after another client reads the data before the commit, that data is then incorrect. The exception is created when the other client attempts to commit the incorrect data.

Retrieving the key that triggered the exception

It might be useful, when troubleshooting such an exception, to retrieve the key corresponding to the entry that triggered the exception. The benefit of the `OptimisticCollisionException` is it contains the `getKey` method, which returns the object representing that key. The following example demonstrates how to retrieve and print the key when catching `OptimisticCollisionException`:

```

try {
...
} catch (OptimisticCollisionException oce) {
    System.out.println(oce.getKey());
}

```

ObjectGridException causes an OptimisticCollisionException

`OptimisticCollisionException` might be the cause of `ObjectGridException` displaying. If this is the case, you can use the following code to determine the exception type and print out the key. The following code uses the `findRootCause` utility method as described in the section below.

```

try {
...
}
catch (ObjectGridException oe) {
    Throwable root = findRootCause( oe );
    if (root instanceof OptimisticCollisionException) {
        OptimisticCollisionException oce = (OptimisticCollisionException)root;
        System.out.println(oce.getKey());
    }
}

```

General exception handling technique

Knowing the root cause of a Throwable object is helpful in isolating the source of an error. The following example demonstrates how an exception handler uses a utility method to find the root cause of the Throwable object.

Example:

```

static public Throwable findRootCause( Throwable t )
{
    // Start with Throwable that occurred as the root.
    Throwable root = t;

    // Follow cause chain until last Throwable in chain is found.
    Throwable cause = root.getCause();
    while ( cause != null )
    {
        root = cause;
        cause = root.getCause();
    }

    // Return last Throwable in the chain as the root cause.
    return root;
}

```

Running parallel business logic on the data grid (DataGrid API)

The DataGrid API provides a simple programming interface to run business logic over all or a subset of the data grid in parallel with where the data is located.

DataGrid APIs and partitioning

With the DataGrid APIs, a client can send requests to one partition, a subset of partitions, or all the partitions in a data grid. The client can specify a list of keys, and WebSphere eXtreme Scale determines the set of partitions that are hosting the keys. The request is then sent to all the partitions in the set in parallel and the client waits for the results. The client can also send requests without specifying keys, therefore, requests are sent to all partitions.

Agents that are deployed to the data grid do not work in client mode. These agents work directly against the primary shard. Working directly against the primary shard results in maximum performance, allowing tens of thousands or more transactions per second because the agent works with the data at full memory speeds. Working directly with the primary shard also means that an agent can only see data that is within that shard. This provides some interesting opportunities that cannot be done on a client.

A typical eXtreme Scale client must be able determine the partition from the transaction, because the client needs to route the request. If an agent is directly attached to a shard, then no routing is needed. All requests go against that shard.

Because the agent is directly attached to a shard, data in other maps in the shard can be accessed without worrying about common partitioning keys, and so on, because no routing occurs.

DataGrid agents and entity-based Maps

A map contains key objects and value objects. The key object is a generated tuple as is the value. An agent is normally provided with the application specified key objects.

The key object is a generated tuple as is the value. An agent is normally provided with the application specified key objects. This will be the key objects used by the application or Tuples if it is an entity Map. An application using Entities will not want to deal with Tuples directly and would prefer to work with the Java objects mapped to the Entity.

Therefore, an Agent class can implement the EntityAgentMixin interface. This forces the class to implement one more method, getClassForEntity(). This returns the entity class to use with the agent on the server side. The keys are converted to this Entity before invoking the process and reduce methods.

This is a different semantic to a non EntityAgentMixin agent where those methods are provided with just the keys. An agent implementing EntityAgentMixin receives the Entity object which includes keys and values in one object.

Note: If the entity does not exist on the server, the keys are the raw Tuple format of the key instead of the managed entity.

DataGrid API example

The DataGrid APIs support two common grid programming patterns: parallel map and parallel reduce.

Parallel Map

The parallel map allows the entries for a set of keys to be processed and returns a result for each entry processed. The application makes a list of keys and receives a Map of key/result pairs after invoking a Map operation. The result is the result of applying a function to the entry of each key. The function is supplied by the application.

MapGridAgent call flow

When the AgentManager.callMapAgent method is invoked with a collection of keys, the MapGridAgent instance is serialized and sent to each primary partition that the keys resolve to. This means that any instance data stored in the agent can be sent to the server. Each primary partition therefore has one instance of the agent. The process method is invoked for each instance one time for each key that resolves to the partition. The result of each process method is then serialized back to the client and returned to the caller in a Map instance, where the result is represented as the value in the map.

When the AgentManager.callMapAgent method is invoked without a collection of keys, the MapGridAgent instance is serialized and sent to every primary partition. This means that any instance data stored in the agent can be sent to the server. Each primary partition therefore has one instance (partition) of the agent. The processAllEntries method is invoked for each partition. The result of each

processAllEntries method is then serialized back to the client and returned to the caller in a Map instance. The following example assumes there is a Person entity with the following shape:

```
import com.ibm.websphere.projector.annotations.Entity;
import com.ibm.websphere.projector.annotations.Id;
@Entity
public class Person
{
    @Id String ssn;
    String firstName;
    String surname;
    int age;
}
```

The application supplied function is written as a class implementing the MapAgentGrid interface. Following is an example agent showing a function to return the age of a Person multiplied by two.

```
public class DoublePersonAgeAgent implements MapGridAgent, EntityAgentMixin
{
    private static final long serialVersionUID = -2006093916067992974L;

    int lowAge;
    int highAge;

    public Object process(Session s, ObjectMap map, Object key)
    {
        Person p = (Person)key;
        return new Integer(p.age * 2);
    }

    public Map processAllEntries(Session s, ObjectMap map)
    {
        EntityManager em = s.getEntityManager();
        Query q = em.createQuery("select p from Person p where p.age > ?1 and p.age < ?2");
        q.setParameter(1, lowAge);
        q.setParameter(2, highAge);
        Iterator iter = q.getResultIterator();
        Map<Person, Integer> rc = new HashMap<Person, Integer>();
        while(iter.hasNext())
        {
            Person p = (Person)iter.next();
            rc.put(p, (Integer)process(s, map, p));
        }
        return rc;
    }

    public Class getClassForEntity()
    {
        return Person.class;
    }
}
```

This shows the Map agent for doubling a Person. Lets look at the process methods first. The first process method is supplied with the Person to work with. It simply returns double the age of that entry. The second process method is called for each partition and finds all Person objects with an age between lowAge and highAge and returns their ages doubled.

```
Session s = grid.getSession();
ObjectMap map = s.getMap("Person");
AgentManager amgr = map.getAgentManager();

DoublePersonAgeAgent agent = new DoublePersonAgeAgent();

// make a list of keys
ArrayList<Person> keyList = new ArrayList<Person>();
Person p = new Person();
p.ssn = "1";
keyList.add(p);
p = new Person ();
p.ssn = "2";
keyList.add(p);
```

```
// get the results for those entries
Map<Tuple, Object> = amgr.callMapAgent(agent, keyList);
```

This shows a client obtaining a Session and a reference to the Person Map. The agent operation is performed against a specific Map. The AgentManager interface is retrieved from that Map. An instance of the agent to invoke is created and any necessary state is added to the object by setting attributes, there are none in this case. A list of keys are then constructed. A Map with the values for person 1 doubled, and the same values for person 2 are returned.

The agent is then invoked for that set of keys. The agents process method is invoked on each partition with some of the specified keys in the grid in parallel. A Map is returned providing the merged results for the specified key. In this case, a Map with the values holding the age for person 1 doubled and the same for person 2 will be returned.

If the key does not exist, the agent will still be invoked. This gives the agent the opportunity to create the map entry. If using an EntityAgentMixin, the key to process will not be the entity, but will instead be the actual Tuple key value for the entity. If the keys are unknown then it's possible to ask all partitions to find Person objects of a certain shape and return their ages doubled. Here is an example:

```
Session s = grid.getSession();
ObjectMap map = s.getMap("Person");
AgentManager amgr = map.getAgentManager();

DoublePersonAgeAgent agent = new DoublePersonAgeAgent();
agent.lowAge = 20;
agent.highAge = 9999;

Map m = amgr.callMapAgent(agent);
```

The previous example shows the AgentManager being obtained for the Person Map, and the agent constructed and initialized with the low and high ages for Persons of interest. The agent is then invoked using the callMapAgent method. Notice, no keys are supplied. This causes the ObjectGrid to invoke the agent on every partition in the grid in parallel and then return the merged results to the client. This will find all Person objects in the grid with an age between low and high and calculate the age of those Person objects doubled. This shows how the grid apis can be used to run a query to find entities matching a certain query. The agent is simply serialized and transported by the ObjectGrid to the partitions with the needed entries. The results are similarly serialized for transport back to the client. Care needs to be taken with the Map APIs. If the ObjectGrid was hosting tera bytes of objects and running on a lot of servers then potentially this would overwhelm anything but the largest machines running the client. This should be used to processing a small subset. If a large subset needs processing then we recommend using a reduce agent to do the processing out in the grid rather than on a client.

Parallel Reduction or aggregation agents

This style of programming processes a subset of the entries and calculates a single result for the group of entries. Examples of such a result would be:

- minimum value
- maximum value
- some other business specific function

A reduce agent is coded and invoked in a very similar manner to the Map agents.

ReduceGridAgent call flow

When the `AgentManager.callReduceAgent` method is invoked with a collection of keys, the `ReduceGridAgent` instance is serialized and sent to each primary partition that the keys resolve to. This means that any instance data stored in the agent can be sent to the server. Each primary partition therefore has one instance of the agent. The `reduce(Session s, ObjectMap map, Collection keys)` method is invoked once per instance (partition) with the subset of keys that resolves to the partition. The result of each reduce method is then serialized back to the client. The `reduceResults` method is invoked on the client `ReduceGridAgent` instance with the collection of each result from each remote reduce invocation. The result from the `reduceResults` method is returned to the caller of the `callReduceAgent` method.

When the `AgentManager.callReduceAgent` method is invoked without a collection of keys, the `ReduceGridAgent` instance is serialized and sent to each primary partition. This means that any instance data stored in the agent can be sent to the server. Each primary partition therefore has one instance of the agent. The `reduce(Session s, ObjectMap map)` method is invoked once per instance (partition). The result of each reduce method is then serialized back to the client. The `reduceResults` method is invoked on the client `ReduceGridAgent` instance with the collection of each result from each remote reduce invocation. The result from the `reduceResults` method is returned to the caller of the `callReduceAgent` method. Here is an example of a reduce agent that simply adds the ages of the matching entries.

```
public class SumAgeReduceAgent implements ReduceGridAgent, EntityAgentMixin
{
    private static final long serialVersionUID = 2521080771723284899L;

    int lowAge;
    int highAge;

    public Object reduce(Session s, ObjectMap map, Collection keyList)
    {
        Iterator<Person> iter = keyList.iterator();
        int sum = 0;
        while (iter.hasNext())
        {
            Person p = iter.next();
            sum += p.age;
        }
        return new Integer(sum);
    }

    public Object reduce(Session s, ObjectMap map)
    {
        EntityManager em = s.getEntityManager ();
        Query q = em.createQuery("select p from Person p where p.age > ?1 and p.age < ?2");
        q.setParameter(1, lowAge);
        q.setParameter(2, highAge);
        Iterator<Person> iter = q.getResultIterator();
        int sum = 0;
        while(iter.hasNext())
        {
            sum += iter.next().age;
        }
        return new Integer(sum);
    }

    public Class getClassForEntity()
    {
        return Person.class;
    }
}
```

The previous example shows the agent. The agent has three important parts. The first allows a specific set of entries to be processed without a query. It simply iterates over the set of entries adding the ages. The sum is returned from the method. The second uses a query to select the entries to be aggregated. It then sums all the matching Person ages. The third method is used to aggregate the

results from each partition to a single result. The ObjectGrid performs the entry aggregation in parallel across the grid. Each partition produces an intermediate result that must be aggregated with other partition intermediate results. This third method performs that task. In the following example the agent is invoked, and the ages of all Persons with ages between 10 and 20 exclusively are aggregated:

```
Session s = grid.getSession();
ObjectMap map = s.getMap("Person");
AgentManager amgr = map.getAgentManager();

SumAgeReduceAgent agent = new SumAgeReduceAgent();

Person p = new Person();
p.ssn = "1";
ArrayList<Person> list = new ArrayList<Person>();
list.add(p);
p = new Person ();
p.ssn = "2";
list.add(p);
Integer v = (Integer)amgr.callReduceAgent(agent, list);
```

Agent functions

The agent is free to do ObjectMap or EntityManager operations within the local shard where it is running. The agent receives a Session and can add, update, query, read, or remove data from the partition the Session represents. Some applications will only query data from the grid, but you can also write an agent to increment all the Person ages by 1 that match a certain query. There is a transaction on the Session when the agent is called, and is committed when the agent returns unless an exception is thrown

Error handling

If a map agent is invoked with an unknown key then the value that is returned is an error object implementing the EntryErrorValue interface.

Transactions

A map agent runs in a separate transaction from the client. Agent invocations may be grouped into a single transaction. If an agent fails (throws an exception), the transaction is rolled-back. Any agents that ran successfully in a transaction will rollback with the failed agent. The AgentManager will rerun the rolled-back agents that ran successfully in a new transaction.

For more information, consult the DataGrid API documentation.

Configuring clients with WebSphere eXtreme Scale

You can configure a WebSphere eXtreme Scale client based on your requirements such as the need to override settings.

Override plug-ins

You can override the following plug-ins on a client:

- **ObjectGrid plug-ins**
 - TransactionCallback plug-in
 - ObjectGridEventListener plug-in
- **BackingMap plug-ins**

- Evictor plug-in
- MapEventListener plug-in
- numberOfBuckets attribute
- ttlEvictorType attribute
- timeToLive attribute

Configure the client with XML

An ObjectGrid XML file can be used to alter settings on the client side. To change the settings on a WebSphere eXtreme Scale client, you must create an ObjectGrid XML file that is similar in structure to the file that was used for the WebSphere eXtreme Scale server.

Assume that the following XML file was paired with a deployment policy XML file, and these files were used to start a WebSphere eXtreme Scale server.

companyGridServerSide.xml

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">

  <objectGrids>
    <objectGrid name="CompanyGrid">
      <bean id="TransactionCallback"
        className="com.company.MyTxCallback" />
      <bean id="ObjectGridEventListener"
        className="com.company.MyOgEventListener" />
      <backingMap name="Customer"
        pluginCollectionRef="customerPlugins" />
      <backingMap name="Item" />
      <backingMap name="OrderLine" numberOfBuckets="1049"
        timeToLive="1600" ttlEvictorType="LAST_ACCESS_TIME" />
      <backingMap name="Order" lockStrategy="PESSIMISTIC"
        pluginCollectionRef="orderPlugins" />
    </objectGrid>
  </objectGrids>

  <backingMapPluginCollections>
    <backingMapPluginCollection id="customerPlugins">
      <bean id="Evictor"
        className="com.ibm.websphere.objectgrid.plugins.builtins.LRUEvictor" />
      <bean id="MapEventListener"
        className="com.company.MyMapEventListener" />
    </backingMapPluginCollection>
    <backingMapPluginCollection id="orderPlugins">
      <bean id="MapIndexPlugin"
        className="com.company.MyMapIndexPlugin" />
    </backingMapPluginCollection>
  </backingMapPluginCollections>
</objectGridConfig>
```

On a WebSphere eXtreme Scale server, the ObjectGrid instance named CompanyGrid behaves as defined by the companyGridServerSide.xml file. By default, the CompanyGrid client has the same settings as the CompanyGrid instance running on the server. However, some of the settings can be overridden on the client, as follows:

1. Create a client-specific ObjectGrid instance.
2. Copy the ObjectGrid XML file that was used to open the server.
3. Edit the new file to customize for the client side.
 - To set or update any of the attributes on the client, specify a new value or change the existing value.
 - To remove a plug-in from the client, use the empty string as the value for the className attribute.

- To change an existing plug-in, specify a new value for the className attribute.
- You can also add any plug-in supported for a client override: TRANSACTION_CALLBACK, OBJECTGRID_EVENT_LISTENER, EVICTOR, MAP_EVENT_LISTENER.

4. Create a client with the newly created client-override XML file.

The following ObjectGrid XML file can be used to specify some of the attributes and plug-ins on the CompanyGrid client.

companyGridClientSide.xml

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">

  <objectGrids>
    <objectGrid name="CompanyGrid">
      <bean id="TransactionCallback"
        className="com.company.MyClientTxCallback" />
      <bean id="ObjectGridEventListener" className="" />
      <backingMap name="Customer" numberOfBuckets="1429"
        pluginCollectionRef="customerPlugins" />
      <backingMap name="Item" />
      <backingMap name="OrderLine" numberOfBuckets="701"
        timeToLive="800" ttlEvictorType="LAST_ACCESS_TIME" />
      <backingMap name="Order" lockStrategy="PESSIMISTIC"
        pluginCollectionRef="orderPlugins" />
    </objectGrid>
  </objectGrids>

  <backingMapPluginCollections>
    <backingMapPluginCollection id="customerPlugins">
      <bean id="Evictor"
        className="com.ibm.websphere.objectgrid.plugins.builtins.LRUEvictor" />
      <bean id="MapEventListener" className="" />
    </backingMapPluginCollection>
    <backingMapPluginCollection id="orderPlugins">
      <bean id="MapIndexPlugin"
        className="com.company.MyMapIndexPlugin" />
    </backingMapPluginCollection>
  </backingMapPluginCollections>
</objectGridConfig>
```

- The TransactionCallback on the client is com.company.MyClientTxCallback instead of the server-side setting of com.company.MyTxCallback.
- The client does not have an ObjectGridEventListener plug-in because the className value is the empty string.
- The client sets the numberOfBuckets to 1429 for the Customer backingMap, retains its Evictor plug-in, and removes the MapEventListener plug-in.
- The numberOfBuckets and timeToLive attributes of the OrderLine backingMap have changed
- Although a different lockStrategy attribute is specified, there is no effect because the lockStrategy attribute is not supported for a client override.

To create the CompanyGrid client using the companyGridClientSide.xml file, pass the ObjectGrid XML file as a URL to one of the connect methods on the ObjectGridManager.

Creating the client for XML

```
ObjectGridManager ogManager =
  ObjectGridManagerFactory.ObjectGridManager();
ClientClusterContext clientClusterContext =
  ogManager.connect("MyServer1.company.com:2809", null, new URL(
    "file:xml/companyGridClientSide.xml"));
```

Configure the client programmatically

You can also override client-side ObjectGrid settings programmatically. Create an ObjectGridConfiguration object that is similar in structure to the server-side ObjectGrid instance. The following code creates a client-side ObjectGrid instance that is functionally equivalent to the client override in the previous section which uses an XML file.

```
client-side override programmatically
ObjectGridConfiguration companyGridConfig = ObjectGridConfigFactory
    .createObjectGridConfiguration("CompanyGrid");
Plugin txCallbackPlugin = ObjectGridConfigFactory.createPlugin(
    PluginType.TRANSACTION_CALLBACK, "com.company.MyClientTxCallback");
companyGridConfig.addPlugin(txCallbackPlugin);

Plugin ogEventListenerPlugin = ObjectGridConfigFactory.createPlugin(
    PluginType.OBJECTGRID_EVENT_LISTENER, "");
companyGridConfig.addPlugin(ogEventListenerPlugin);

BackingMapConfiguration customerMapConfig = ObjectGridConfigFactory
    .createBackingMapConfiguration("Customer");
customerMapConfig.setNumberOfBuckets(1429);
Plugin evictorPlugin = ObjectGridConfigFactory.createPlugin(PluginType.EVICTOR,
    "com.ibm.websphere.objectgrid.plugins.builtins.LRUEvictor");
customerMapConfig.addPlugin(evictorPlugin);

companyGridConfig.addBackingMapConfiguration(customerMapConfig);

BackingMapConfiguration orderLineMapConfig = ObjectGridConfigFactory
    .createBackingMapConfiguration("OrderLine");
orderLineMapConfig.setNumberOfBuckets(701);
orderLineMapConfig.setTimeToLive(800);
orderLineMapConfig.setTtlEvictorType(TTLType.LAST_ACCESS_TIME);

companyGridConfig.addBackingMapConfiguration(orderLineMapConfig);

List ogConfigs = new ArrayList();
ogConfigs.add(companyGridConfig);

Map overrideMap = new HashMap();
overrideMap.put(CatalogServerProperties.DEFAULT_DOMAIN, ogConfigs);

ogManager.setOverrideObjectGridConfigurations(overrideMap);
ClientClusterContext client = ogManager.connect(catalogServerAddresses, null, null);
ObjectGrid companyGrid = ogManager.getObjectGrid(client, objectGridName);
```

The ogManager instance of the ObjectGridManager interface checks for overrides only in the ObjectGridConfiguration and BackingMapConfiguration objects that you include in the overrideMap Map. For instance, the previous code overrides the number of buckets on the OrderLine Map. However, the Order map remains unchanged on the client side because no configuration for that map is included.

Configure the client in the Spring Framework

Client-side ObjectGrid settings can also be overridden using the Spring Framework. The following example XML file shows how to build an ObjectGridConfiguration element, and use it to override some client side settings. This example calls the same APIs that are demonstrated in the programmatic configuration. The example is also functionally equivalent to the example in the ObjectGrid XML configuration.

```
client configuration with Spring
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE beans PUBLIC "-//SPRING//DTD BEAN//EN"
    "http://www.springframework.org/dtd/spring-beans.dtd">
<beans>
  <bean id="companyGrid" factory-bean="manager" factory-method="getObjectGrid"
    singleton="true">
    <constructor-arg type="com.ibm.websphere.objectgrid.ClientClusterContext">
      <ref bean="client" />
    </constructor-arg>
    <constructor-arg type="java.lang.String" value="CompanyGrid" />
  </bean>
</beans>
```

```

</bean>

<bean id="manager" class="com.ibm.websphere.objectgrid.ObjectGridManagerFactory"
  factory-method="getObjectGridManager" singleton="true">
  <property name="overrideObjectGridConfigurations">
    <map>
      <entry key="DefaultDomain">
        <list>
          <ref bean="ogConfig" />
        </list>
      </entry>
    </map>
  </property>
</bean>

<bean id="ogConfig"
  class="com.ibm.websphere.objectgrid.config.ObjectGridConfigFactory"
  factory-method="createObjectGridConfiguration">
  <constructor-arg type="java.lang.String">
    <value>CompanyGrid</value>
  </constructor-arg>
  <property name="plugins">
    <list>
      <bean class="com.ibm.websphere.objectgrid.config.ObjectGridConfigFactory"
        factory-method="createPlugin">
        <constructor-arg type="com.ibm.websphere.objectgrid.config.PluginType"
          value="TRANSACTION_CALLBACK" />
        <constructor-arg type="java.lang.String"
          value="com.company.MyClientTxCallback" />
        </bean>
      <bean class="com.ibm.websphere.objectgrid.config.ObjectGridConfigFactory"
        factory-method="createPlugin">
        <constructor-arg type="com.ibm.websphere.objectgrid.config.PluginType"
          value="OBJECTGRID_EVENT_LISTENER" />
        <constructor-arg type="java.lang.String" value="" />
        </bean>
      </list>
    </property>
    <property name="backingMapConfigurations">
      <list>
        <bean class="com.ibm.websphere.objectgrid.config.ObjectGridConfigFactory"
          factory-method="createBackingMapConfiguration">
          <constructor-arg type="java.lang.String" value="Customer" />
          <property name="plugins">
            <bean class="com.ibm.websphere.objectgrid.config.ObjectGridConfigFactory"
              factory-method="createPlugin">
              <constructor-arg type="com.ibm.websphere.objectgrid.config.PluginType"
                value="EVICTOR" />
              <constructor-arg type="java.lang.String"
                value="com.ibm.websphere.objectgrid.plugins.builtins.LRUEvictor" />
              </bean>
            </property>
            <property name="numberOfBuckets" value="1429" />
          </bean>
          <bean class="com.ibm.websphere.objectgrid.config.ObjectGridConfigFactory"
            factory-method="createBackingMapConfiguration">
            <constructor-arg type="java.lang.String" value="OrderLine" />
            <property name="numberOfBuckets" value="701" />
          </bean>
        </list>
      </property>
      <property name="timeToLive" value="800" />
      <property name="ttlEvictorType">
        <value type="com.ibm.websphere.objectgrid.
          TTLType">LAST_ACCESS_TIME</value>
      </property>
    </bean>
  </list>
</property>
</bean>

  <bean id="client" factory-bean="manager" factory-method="connect"
    singleton="true">
    <constructor-arg type="java.lang.String">
      <value>localhost:2809</value>
    </constructor-arg>
    <constructor-arg
      type="com.ibm.websphere.objectgrid.security.
        config.ClientSecurityConfiguration">
      <null />
    </constructor-arg>
    <constructor-arg type="java.net.URL">

```

```
<null />
  </constructor-arg>
</bean>
</beans>
```

After creating the XML file, load the file and build the ObjectGrid with the following code snippet.

```
BeanFactory beanFactory = new XmlBeanFactory(new
UrlResource("file:test/companyGridSpring.xml"));

ObjectGrid companyGrid = (ObjectGrid) beanFactory.getBean("companyGrid");
```

Read about the Spring framework integration overview for more information on creating an XML descriptor file.

Disable the client near cache

The near cache is enabled by default when locking is configured as optimistic or none. Clients do not maintain a near cache when the locking setting is configured as pessimistic. To disable the near cache, you must set the `numberOfBuckets` attribute to 0 in the client override ObjectGrid descriptor file.

Tracking map updates by an application

When an application is making changes to a Map during a transaction, a LogSequence object tracks those changes. If the application changes an entry in the map, a corresponding LogElement object provides the details of the change.

Loaders are given a LogSequence object for a particular map whenever an application calls for a flush or commit to the transaction. The Loader iterates over the LogElement objects within the LogSequence object and applies each LogElement object to the backend.

ObjectGridEventListener listeners that are registered with an ObjectGrid also use LogSequence objects. These listeners are given a LogSequence object for each map in a committed transaction. Applications can use these listeners to wait for certain entries to change, like a trigger in a conventional database.

The following log-related interfaces or classes are provided by the eXtreme Scale framework:

- `com.ibm.websphere.objectgrid.plugins.LogElement`
- `com.ibm.websphere.objectgrid.plugins.LogSequence`
- `com.ibm.websphere.objectgrid.plugins.LogSequenceFilter`
- `com.ibm.websphere.objectgrid.plugins.LogSequenceTransformer`

LogElement interface

A LogElement represents an operation on an entry during a transaction. A LogElement object has several methods to get its various attributes. The most commonly used attributes are the type and the current value attributes fetched by `getType()` and `getCurrentValue()`.

The type is represented by one of the constants defined in the LogElement interface: INSERT, UPDATE, DELETE, EVICT, FETCH, or TOUCH.

The current value represents the new value for the operation if it is INSERT, UPDATE or FETCH. If the operation is TOUCH, DELETE, or EVICT, then the current value is null. This value can be cast to ValueProxyInfo when a ValueInterface is in use.

See the API documentation for more details on the LogElement interface.

LogSequence interface

In most transactions, operations to more than one entry in a map occur, so multiple LogElement objects are created. You should create an object that behaves as a composite of multiple LogElement objects. The LogSequence interface serves this purpose by containing a list of LogElement objects.

See the API documentation for more details on the LogSequence interface.

Using LogElement and LogSequence

LogElement and LogSequence are widely used in eXtreme Scale and by ObjectGrid plug-ins that are written by users when operations are propagated from one component or server to another component or server. For example, a LogSequence object can be used by the distributed ObjectGrid transaction propagation function to propagate the changes to other servers, or it can be applied to the persistence store by the loader. LogSequence is mainly used by the following interfaces.

- com.ibm.websphere.objectgrid.plugins.ObjectGridEventListener
- com.ibm.websphere.objectgrid.plugins.Loader
- com.ibm.websphere.objectgrid.plugins.Evictor
- com.ibm.websphere.objectgrid.Session

Loader example

This section demonstrates how the LogSequence and LogElement objects are used in a Loader. A Loader is used to load data from and persist data into a persistent store. The batchUpdate method of the Loader interface uses LogSequence object:

```
void batchUpdate(TxID txid, LogSequence sequence) throws  
    LoaderException, OptimisticCollisionException;
```

The batchUpdate method is called when an ObjectGrid needs to apply all current changes to the Loader. The Loader is given a list of LogElement objects for the map, encapsulated in a LogSequence object. The implementation of the batchUpdate method must iterate over the changes and apply them to the backend. The following code snippet demonstrates how a Loader uses a LogSequence object. The snippet iterates over the set of changes and builds up three batch Java database connectivity (JDBC) statements: inserts, updates, and deletes:

```
public void batchUpdate(TxID tx, LogSequence sequence) throws LoaderException  
{  
    // Get a SQL connection to use.  
    Connection conn = getConnection(tx);  
    try  
    {  
        // Process the list of changes and build a set of prepared  
        // statements for executing a batch update, insert, or delete  
        // SQL operations. The statements are cached in stmtCache.  
        Iterator iter = sequence.getPendingChanges();  
        while ( iter.hasNext() )  
        {
```



```

        LogElement logElement = (LogElement)iter.next();
        Object key = logElement.getCacheEntry().getKey();
        Object value = logElement.getCurrentValue();
        switch ( logElement.getType().getCode() )
        {
            case LogElement.CODE_INSERT:
                buildBatchSQLInsert( key, value, conn );
                break;
            case LogElement.CODE_UPDATE:
                buildBatchSQLUpdate( key, value, conn );
                break;
            case LogElement.CODE_DELETE:
                buildBatchSQLDelete( key, conn );
                break;
        }
    }
    // Run the batch statements that were built by above loop.
    Collection statements = getPreparedStatementCollection( tx, conn );
    iter = statements.iterator();
    while ( iter.hasNext() )
    {
        PreparedStatement pstmt = (PreparedStatement) iter.next();
        pstmt.executeBatch();
    }
} catch (SQLException e)
{
    LoaderException ex = new LoaderException(e);
    throw ex;
}
}
}

```

The previous sample illustrates the high-level logic of processing the `LogSequence` argument. However, the sample does not illustrate the details of how an SQL insert, update, or delete statement is built. The `getPendingChanges` method is called on the `LogSequence` argument to obtain an iterator of `LogElement` objects that a `Loader` needs to process, and the `LogElement.getType().getCode()` method is used to determine whether a `LogElement` is for an SQL insert, update, or delete operation.

Evictor sample

You can also use `LogSequence` and `LogElement` objects with an `Evictor`. An `Evictor` is used to evict the map entries from the backing map based on certain criteria. The `apply` method of the `Evictor` interface uses `LogSequence`.

```

/**
 * This is called during cache commit to allow the evictor to track object usage
 * in a backing map. This will also report any entries that have been successfully
 * evicted.
 *
 * @param sequence LogSequence of changes to the map
 */
void apply(LogSequence sequence);

```

LogSequenceFilter and LogSequenceTransformer interfaces

Sometimes, it is necessary to filter the `LogElement` objects so that only `LogElement` objects with certain criteria are accepted, and reject other objects. For example, you might want to serialize a certain `LogElement` based on some criterion.

`LogSequenceFilter` solves this problem with the following method.

```

public boolean accept (LogElement logElement);

```

This method returns true if the given LogElement should be used in the operation, and returns false if the given LogElement should not be used.

LogSequenceTransformer is a class that uses the LogSequenceFilter function. It uses the LogSequenceFilter to filter out some LogElement objects and then serialize the accepted LogElement objects. This class has two methods. The first method follows.

```
public static void serialize(Collection logSequences, ObjectOutputStream stream,
    LogSequenceFilter filter, DistributionMode mode) throws IOException
```

This method allows the caller to provide a filter for determining which LogElements to include in the serialization process. The DistributionMode parameter allows the caller to control the serialization process. For example, if the distribution mode is invalidation only, then there is no need to serialize the value. The second method of this class is the inflate method, as follows.

```
public static Collection inflate(ObjectInputStream stream, ObjectGrid
    objectGrid) throws IOException, ClassNotFoundException
```

The inflate method reads the log sequence serialized form, which was created by the serialize method, from the provided object input stream.

Enabling client-side map replication

You can also enable replication of maps on the client side to make data available faster.

With eXtreme Scale, you can replicate a server map to one or more clients by using asynchronous replication. A client can request a local read-only copy of a server side map by using the ClientReplicableMap.enableClientReplication method.

```
void enableClientReplication(Mode mode, int[] partitions,
    ReplicationMapListener listener) throws ObjectGridException;
```

The first parameter is the replication mode. This mode can be a continuous replication or a snapshot replication. The second parameter is an array of partition IDs that represent the partitions from which to replicate the data. If the value is null or an empty array, the data is replicated from all the partitions. The last parameter is a listener to receive client replication events. See ClientReplicableMap and ReplicationMapListener in the API documentation for details.

After the replication is enabled, then the server starts to replicate the map to the client. The client is eventually only a few transactions behind the server at any point in time.

Chapter 3. Accessing data with the REST data service

Develop applications that perform operations using REST data service protocols.

`startOgServer`

Operations with the REST data service

After you start the eXtreme Scale REST data service, you can use any HTTP client to interact with it. A Web browser, PHP client, Java client or WCF Data Services client can be used to issue any of the supported request operations.

The REST service implements a subset of the Microsoft Atom Publishing Protocol: Data Services URI and Payload Extensions specification, Version 1.0 which is part of OData protocol. This topic describes which of the features of the specification are supported and how they are mapped to eXtreme Scale.

Service root URI

Microsoft WCF Data Services typically defines a service per data source or entity model. The eXtreme Scale REST data service defines a service per defined ObjectGrid. Each ObjectGrid that is defined in the eXtreme Scale ObjectGrid client override XML file is automatically exposed as a separate REST service root.

The URI for the service root is:

`http://host:port/contextroot/restservice/gridname`

Where:

- *contextroot* is defined when you deploy the REST data service application, and depends on the application server
- *gridname* is the name of the ObjectGrid

Request types

The following list describes the Microsoft WCF Data Services request types which the eXtreme Scale REST data service supports. For details about each request type that WCF Data Services supports, see: MSDN: Request Types.

Insert request types

Clients can insert resources using the POST HTTP verb with the following limitations:

- InsertEntity Request: Supported.
- InsertLink request: Supported.
- InsertMediaResource request: Not supported due to media resource support restriction.

For additional information, see: MSDN: Insert Request Types.

Update request types

Clients can update resources using the PUT and MERGE HTTP verbs with the following limitations:

- UpdateEntity Request: Supported.

- UpdateComplexType Request: Not Supported due to complex type restriction.
- UpdatePrimitiveProperty Request: Supported.
- UpdateValue Request: Supported.
- UpdateLink Request: Supported.
- UpdateMediaResource Request: Not supported due to media resource support restriction.

For additional information, see: MSDN: Insert Request types.

Delete request types

Clients can delete resources using the DELETE HTTP verb with the following limitations:

- DeleteEntity Request: Supported.
- DeleteLink Request: Supported.
- DeleteValue request: Supported.

For additional information, see: MSDN: Delete Request Types.

Retrieve request types

Clients can retrieve resources using the GET HTTP verb with the following limitations:

- RetrieveEntitySet Request: Supported.
- RetrieveEntity Request: Supported.
- RetrieveComplexType Request: Not supported due to complex type restriction.
- RetrievePrimitiveProperty Request: Supported.
- RetrieveValue Request: Supported.
- RetrieveServiceMetadata Request: Supported.
- RetrieveServiceDocument Request: Supported.
- RetrieveLink Request: Supported.
- Retrieve Request Containing a Customizable Feed Mapping: Not supported
- RetrieveMediaResource: Not supported due to media resource restriction.

For additional information, see: MSDN: Retrieve Request Types.

System query options

Queries are supported which allow clients to identify a collection of entities or a single entity. System query options are specified in a data service URI and are supported with the following limitations:

- \$expand: Supported
- \$filter: Supported.
- \$orderby: Supported.
- \$format: Not supported. The acceptable format is identified in the HTTP Accept request header.
- \$skip: Supported
- \$top: Supported

For additional information, see: MSDN: System Query Options.

Partition routing

Partition routing is based on the root entity. A request URI infers a root entity if its resource path starts with a root entity or with an entity that has a direct or indirect association to the entity. In a partitioned environment, any request that cannot infer a root entity will be rejected. Any request that infers a root entity will be routed to the correct partition.

For additional information on defining a schema with associations and root entities, see Scalable data model in eXtreme Scale and Partitioning.

Invoke request

Invoke requests are not supported. For additional information, see MSDN: Invoke Request.

Batch request

Clients can batch multiple Change Sets or Query Operations within a single request. This can reduce the number of round trips to the server and allows multiple requests to participate in a single transaction. For additional information, see MSDN: Batch Request.

Tunneled requests

Tunneled requests are not supported. For additional information, see MSDN: Tunneled Requests.

Request protocols for the REST data service

In general, the protocols for interacting with the REST service are the same as those described in the WCF Data Services AtomPub protocol. However, eXtreme Scale does provide additional details, from eXtreme Scale Entity Model perspective. Users are expected to be familiar with the WCF Data Services protocols before reading this section. Alternatively, users can read this section with the WCF Data Services protocol section.

Examples are provided to illustrate the request and response. These examples apply to both the eXtreme Scale REST data service and WCF Data Services. Because Web browsers can only retrieve data, the CRUD (create, update and delete) operations must be performed by another client such as Java, JavaScript, RUBY or PHP.

Retrieve requests with the REST data service

A RetrieveEntity Request is used by a client to retrieve an eXtreme Scale entity. The response payload contains the entity data in AtomPub or JSON format. Also, the system operator \$expand can be used to expand the relations. The relations are represented in line within the data service response as an Atom Feed Document, which is a to-many relation, or an Atom Entry Document which is a to-one relation.

Tip: For more details on the RetrieveEntity protocol defined in WCF Data Services, refer to MSDN: RetrieveEntity Request.

Retrieving an entity

The following RetrieveEntity example retrieves a Customer entity with key.

AtomPub

- Method
GET
- Request URI:
`http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Customer('ACME')`
- Request Header:
Accept: application/atom+xml
- Request Payload:
None
- Response Header:
Content-Type: application/atom+xml
- Response Payload:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<entry xml:base = "http://localhost:8080/wxsrestservice/
restservice" xmlns:d= "http://schemas.microsoft.com/ado/2007/
08/dataservices" xmlns:m = "http://schemas.microsoft.com/ado/2007/
08/dataservices/metadata" xmlns = "http://www.w3.org/2005/Atom">

<category term = "NorthwindGridModel.Customer" scheme = "http://
schemas.microsoft.com/ado/2007/08/dataservices/scheme"/>
<id>http://localhost:8080/wxsrestservice/restservice/
NorthwindGrid/Customer('ACME')</id>
<title type = "text"/>
<updated>2009-12-16T19:52:10.593Z</updated>
<author>
<name/>
</author>
<link rel = "edit" title = "Customer" href = "Customer(
'ACME')"/>
<link rel = "http://schemas.microsoft.com/ado/2007/08/
dataservices/related/
orders" type = "application/atom+xml;type=feed" title =
"orders" href = "Customer('ACME')/orders"/>
<content type = "application/xml">
<m:properties>
<d:customerId>ACME</d:customerId>
<d:city m:null = "true"/>
<d:companyName>RoaderRunner</d:companyName>
<d:contactName>ACME</d:contactName>
<d:country m:null = "true"/>
<d:version m:type = "Edm.Int32">3</d:version>
</m:properties>
</content>
</entry>
```
- Response Code:
200 OK

JSON

- Method
GET
- Request URI:

http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/
Customer('ACME')

- Request Header:
Accept: application/json
- Request Payload:
None
- Response Header:
Content-Type: application/json
- Response Payload:

```
{ "d": { "__metadata": { "uri": "http://localhost:8080/wxsrestservice/  
restservice/NorthwindGrid/Customer('ACME')",  
"type": "NorthwindGridModel.Customer",  
"customerId": "ACME",  
"city": null,  
"companyName": "RoaderRunner",  
"contactName": "ACME",  
"country": null,  
"version": 3,  
"orders": { "__deferred": { "uri": "http://localhost:8080/  
wxsrestservice/restservice/  
NorthwindGrid/Customer('ACME')/orders" } } } }
```
- Response Code:
200 OK

Queries

A query can also be used with a RetrieveEntitySet or RetrieveEntity request. A query is specified by the system \$filter operator.

For details on the \$filter operator, refer to: MSDN: Filter System Query Option (\$filter)

The OData protocol supports several common expressions. The eXtreme Scale REST data service supports a subset of the expressions defined in the specification:

- Boolean expressions:
 - eq, ne, lt, le, gt, ge
 - negate
 - not
 - parenthesis
 - and, or
- Arithmetic expressions:
 - add
 - sub
 - mul
 - div
- Primitive literals
 - String
 - date-time
 - decimal
 - single
 - double

- int16
- int32
- int64
- binary
- null
- byte

The following expressions are *not* available:

- Boolean expressions:
 - isof
 - cast
- Method call expressions
- Arithmetic expressions:
 - mod
- Primitive literals:
 - Guid
- Member expressions

For a complete list and description of the expressions that are available in Microsoft WCF Data Services, see section 2.2.3.6.1.1 : Common Expression Syntax.

The following example demonstrates a RetrieveEntity request with a query. In this example, all customers whose contact name is "RoadRunner" are retrieved. The only customer which matches this filter is Customer('ACME') as shown in the response payload.

Restriction: This query will only work for non-partitioned entities. If Customer is partitioned, then the key belonging to the customer is required.

AtomPub

- Method: GET
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/ Customer?$filter=contactName eq 'RoadRunner'`
- Request Header: Accept: application/atom+xml
- Input Payload: None
- Response Header: Content-Type: application/atom+xml
- Response Payload:

```
<?xml version="1.0" encoding="iso-8859-1"?>
<feed
  xml:base="http://localhost:8080/wxsrestservice/restservice"
  xmlns:d="http://schemas.microsoft.com/ado/2007/08/
    dataservices"
  xmlns:m="http://schemas.microsoft.com/ado/2007/08/
    dataservices/metadata"
  xmlns="http://www.w3.org/2005/Atom">
  <title type="text">Customer</title>
  <id> http://localhost:8080/wxsrestservice/restservice/
    NorthwindGrid/Customer </id>
  <updated>2009-09-16T04:59:28.656Z</updated>
  <link rel="self" title="Customer" href="Customer" />
  <entry>
    <category term="NorthwindGridModel.Customer"
      scheme="http://schemas.microsoft.com/ado/2007/08/
```



```

    dataservices/scheme" />
    <id>
http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/
Customer('ACME')</id>
    <title type="text" />
    <updated>2009-09-16T04:59:28.656Z</updated>
    <author>
    <name />
    </author>
    <link rel="edit" title="Customer" href="Customer('ACME')" />
    <link
    rel="http://schemas.microsoft.com/ado/2007/08/dataservices/
    related/orders"
    type="application/atom+xml;type=feed" title="orders"
    href="Customer('ACME')/orders" />
    <content type="application/xml">
    <m:properties>
    <d:customerId>ACME</d:customerId>
    <d:city m:null = "true"/>
    <d:companyName>RoaderRunner</d:companyName>
    <d:contactName>ACME</d:contactName>
    <d:country m:null = "true"/>
    <d:version m:type = "Edm.Int32">3</d:version>
    </m:properties>
    </content>
    </entry>
</feed>

```

- Response Code: 200 OK

JSON

- Method: GET
- Request URI:


```
http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/
Customer?$filter=contactName eq 'RoadRunner'
```
- Request Header: Accept: application/json
- Request Payload: None
- Response Header: Content-Type: application/json
- Response Payload:


```
{
  "d": [
    {
      "__metadata": {
        "uri": "http://localhost:8080/wxsrestservice/
restservice/NorthwindGrid/Customer('ACME')",
        "type": "NorthwindGridModel.Customer",
        "customerId": "ACME",
        "city": null,
        "companyName": "RoaderRunner",
        "contactName": "ACME",
        "country": null,
        "version": 3,
        "orders": {
          "__deferred": {
            "uri": "http://localhost:8080/
wxsrestservice/restservice/NorthwindGrid/
Customer('ACME')/orders"
          }
        }
      }
    }
  ]
}
```
- Response Code: 200 OK

System operator \$expand

The system operator \$expand can be used to expand associations. The associations are represented in line in the data service response. Multi-valued (to-many) associations are represented as an Atom Feed Document or JSON array. Single-valued (to-one) associations, are represented as n Atom Entry Document or JSON object.

For more details on the \$expand system operator, refer to Expand System Query Option (\$expand).

Here is an example of using the \$expand system operator. In this example, we retrieve the entity Customer('IBM') which has an Orders 5000, 5001 and others associated with it. The \$expand clause is set to "orders", so the order collection is expand as inline in the response payload. Only orders 5000 and 5001 are displayed here.

AtomPub

- Method: GET
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Customer('IBM')?$expand=orders`
- Request Header: Accept: application/atom+xml
- Request Payload: None
- Response Header: Content-Type: application/atom+xml
- Response Payload:

```
<?xml version="1.0" encoding="utf-8"?>
<entry xml:base = "http://localhost:8080/wxsrestservice/restservice"
  xmlns:d = "http://schemas.microsoft.com/ado/2007/08/dataservices"
  xmlns:m = "http://schemas.microsoft.com/ado/2007/08/dataservices/
  metadata" xmlns = "http://www.w3.org/2005/Atom">
<category term = "NorthwindGridModel.Customer" scheme = "http://schemas.
microsoft.com/ado/2007/08/dataservices/scheme"/>
  <id>http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/
  Customer('IBM')</id>
  <title type = "text"/>
  <updated>2009-12-16T22:50:18.156Z</updated>
  <author>
    <name/>
  </author><link rel = "edit" title = "Customer" href =
  "Customer('IBM')"/>
  <link rel = "http://schemas.microsoft.com/ado/2007/08/dataservices/
  related/orders" type = "application/atom+xml;type=feed" title =
  "orders" href = "Customer('IBM')/orders">
    <m:inline>
      <feed>
        <title type = "text">orders</title>
        <id>http://localhost:8080/wxsrestservice/restservice/
  NorthwindGrid/Customer('IBM')/orders</id>
        <updated>2009-12-16T22:50:18.156Z</updated>
        <link rel = "self" title = "orders" href = "Customer
  ('IBM')/orders"/>
        <entry>
          <category term = "NorthwindGridModel.Order" scheme =
  "http://schemas.microsoft.com/ado/2007/08/
  dataservices/scheme"/>
          <id>http://localhost:8080/wxsrestservice/restservice/
  NorthwindGrid/Order(orderId=5000,customer_customerId=
  'IBM')</id>
          <title type = "text"/>
          <updated>2009-12-16T22:50:18.156Z</updated>
          <author>
            <name/>
          </author>
          <link rel = "edit" title = "Order" href =
  "Order(orderId=5000,customer_customerId='IBM')"/>
          <link rel = "http://schemas.microsoft.com/ado/2007/08/
  dataservices/related/customer" type = "application/
  atom+xml;type=entry" title = "customer" href =
  "Order(orderId=5000,customer_customerId='IBM')/customer"/>
```

```

    <link rel = "http://schemas.microsoft.com/ado/2007/08/
dataservices/related/orderDetails" type = "application/
atom+xml;type=feed" title = "orderDetails" href =
"Order(orderId=5000,customer_customerId='IBM')/orderDetails"/>
    <content type = "application/xml">
      <m:properties>
        <d:orderId m:type = "Edm.Int32">5000</d:orderId>
        <d:customer_customerId>IBM</d:customer_customerId>
        <d:orderDate m:type = "Edm.DateTime">
2009-12-16T19:46:29.562</d:orderDate>
        <d:shipCity>Rochester</d:shipCity>
        <d:shipCountry m:null = "true"/>
        <d:version m:type = "Edm.Int32">0</d:version>
      </m:properties>
    </content>
  </entry>
  <entry>
    <category term = "NorthwindGridModel.Order" scheme =
"http://schemas.microsoft.com/ado/2007/08/
dataservices/scheme"/>
    <id>http://localhost:8080/wxsrestservice/restservice/
NorthwindGrid/Order(orderId=5001,customer_customerId=
'IBM')</id>
    <title type = "text"/>
    <updated>2009-12-16T22:50:18.156Z</updated>
    <author>
      <name/></author>
    <link rel = "edit" title = "Order" href = "Order(
orderId=5001,customer_customerId='IBM')"/>
    <link rel = "http://schemas.microsoft.com/ado/2007/
08/dataservices/related/customer" type =
"application/atom+xml;type=entry" title =
"customer" href = "Order(orderId=5001,customer_customerId=
'IBM')/customer"/>
    <link rel = "http://schemas.microsoft.com/ado/2007/08/
dataservices/related/orderDetails" type =
"application/atom+xml;type=feed" title =
"orderDetails" href = "Order(orderId=5001,
customer_customerId='IBM')/orderDetails"/>
    <content type = "application/xml">
      <m:properties>
        <d:orderId m:type = "Edm.Int32">5001</d:orderId>
        <d:customer_customerId>IBM</d:customer_customerId>
        <d:orderDate m:type = "Edm.DateTime">2009-12-16T19:
50:11.125</d:orderDate>
        <d:shipCity>Rochester</d:shipCity>
        <d:shipCountry m:null = "true"/>
        <d:version m:type = "Edm.Int32">0</d:version>
      </m:properties>
    </content>
  </entry>
</feed>
</m:inline>
</link>
<content type = "application/xml">
  <m:properties>
    <d:customerId>IBM</d:customerId>
    <d:city m:null = "true"/>
    <d:companyName>IBM Corporation</d:companyName>
    <d:contactName>John Doe</d:contactName>
    <d:country m:null = "true"/>
    <d:version m:type = "Edm.Int32">4</d:version>
  </m:properties>
</content>
</entry>

```

- Response Code: 200 OK

JSON

- Method: GET
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/ Customer('IBM')?$expand=orders`
- Request Header: Accept: `application/json`
- Request Payload: None
- Response Header: Content-Type: `application/json`
- Response Payload:

```
{ "d": { "__metadata": { "uri": "http://localhost:8080/wxsrestservice/
  restservice/NorthwindGrid/Customer('IBM')",
  "type": "NorthwindGridModel.Customer" },
  "customerId": "IBM",
  "city": null,
  "companyName": "IBM Corporation",
  "contactName": "John Doe",
  "country": null,
  "version": 4,
  "orders": [ { "__metadata": { "uri": "http://localhost:8080/
    wxsrestservice/restservice/NorthwindGrid/Order(
      orderId=5000,customer_customerId='IBM')",
    "type": "NorthwindGridModel.Order" },
    "orderId": 5000,
    "customer_customerId": "IBM",
    "orderDate": "\\Date(1260992789562)\\/",
    "shipCity": "Rochester",
    "shipCountry": null,
    "version": 0,
    "customer": { "__deferred": { "uri": "http://localhost:8080/
      wxsrestservice/restservice/NorthwindGrid/Order(
        orderId=5000,customer_customerId='IBM')/customer" } },
    "orderDetails": { "__deferred": { "uri": "http://localhost:
      8080/wxsrestservice/restservice/NorthwindGrid/
      Order(orderId=5000,customer_customerId='IBM')/
      orderDetails" } } },
    { "__metadata": { "uri": "http://localhost:8080/wxsrestservice/
      restservice/NorthwindGrid/Order(orderId=5001,
        customer_customerId='IBM')", "type":
        "NorthwindGridModel.Order" },
    "orderId": 5001,
    "customer_customerId": "IBM",
    "orderDate": "\\Date(1260993011125)\\/",
    "shipCity": "Rochester",
    "shipCountry": null,
    "version": 0,
    "customer": { "__deferred": { "uri": "http://localhost:
      8080/wxsrestservice/restservice/
      NorthwindGrid/Order(orderId=5001,customer_customerId
        ='IBM')/customer" } },
    "orderDetails": { "__deferred": { "uri": "http://localhost:8080/
      wxsrestservice/restservice/NorthwindGrid/Order(
        orderId=5001, customer_customerId='IBM')/
        orderDetails" } } } ] }
```

- Response Code: 200 OK

Retrieving non-entities with REST data services

The REST data service allows you to retrieve more than only entities, such as entity collections and properties.

Retrieve an entity collection

A RetrieveEntitySet Request can be used by a client to retrieve a set of eXtreme Scale entities. The entities are represented as an Atom Feed Document or JSON array in the response payload. For more details on the RetrieveEntitySet protocol defined in WCF Data Services, refer to MSDN: RetrieveEntitySet Request.

The following RetrieveEntitySet request example retrieves all the Order entities associated with the Customer('IBM') entity. Only orders 5000 and 5001 are displayed here.

AtomPub

- Method: GET
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Customer('IBM')/orders`
- Request Header: Accept: application/atom+xml
- Request Payload: None
- Response Header: Content-Type: application/atom+xml
- Response Payload:

```
<?xml version="1.0" encoding="utf-8"?>
<feed xml:base = "http://localhost:8080/wxsrestservice/restservice"
      xmlns:d = "http://schemas.microsoft.com/ado/2007/08/dataservices"
      xmlns:m = "http://schemas.microsoft.com/ado/2007/08/dataservices/
      metadata" xmlns = "http://www.w3.org/2005/Atom">
  <title type = "text">Order</title>
  <id>http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/
  Order</id>
  <updated>2009-12-16T22:53:09.062Z</updated>
  <link rel = "self" title = "Order" href = "Order"/>
  <entry>
    <category term = "NorthwindGridModel.Order" scheme = "http://
    schemas.microsoft.com/
    ado/2007/08/dataservices/scheme"/>
    <id>http://localhost:8080/wxsrestservice/restservice/
    NorthwindGrid/Order(orderId=5000,customer_customerId=
    'IBM')</id>
    <title type = "text"/>
    <updated>2009-12-16T22:53:09.062Z</updated>
    <author>
      <name/>
    </author>
    <link rel = "edit" title = "Order" href = "Order(orderId=5000,
    customer_customerId='IBM')"/>
    <link rel = "http://schemas.microsoft.com/ado/2007/08/
    dataservices/related/customer"
    type = "application/atom+xml;type=entry"
    title = "customer" href = "Order(orderId=5000,
    customer_customerId='IBM')/customer"/>
    <link rel = "http://schemas.microsoft.com/ado/2007/08/
    dataservices/related/orderDetails"
    type = "application/atom+xml;type=feed"
    title = "orderDetails" href = "Order(orderId=5000,
    customer_customerId='IBM')/
    orderDetails"/>
    <content type = "application/xml">
      <m:properties>
        <d:orderId m:type = "Edm.Int32">5000</d:orderId>
        <d:customer_customerId>IBM</d:customer_customerId>
        <d:orderDate m:type = "Edm.DateTime">2009-12-16T19:
        46:29.562</d:orderDate>
        <d:shipCity>Rochester</d:shipCity>
      </m:properties>
    </content>
  </entry>
</feed>
```

```

        <d:shipCountry m:null = "true"/>
        <d:version m:type = "Edm.Int32">0</d:version>
    </m:properties>
</content>
</entry>
<entry>
    <category term = "NorthwindGridModel.Order" scheme = "http://
schemas.microsoft.com/ado/2007/08/dataservices/scheme"/>
    <id>http://localhost:8080/wxsrestservice/restservice/
NorthwindGrid/Order(orderId=5001, customer_customerId='IBM')
</id>
    <title type = "text"/>
    <updated>2009-12-16T22:53:09.062Z</updated>
    <author>
        <name/>
    </author>
    <link rel = "edit" title = "Order" href = "Order(orderId=5001,
customer_customerId='IBM')"/>
    <link rel = "http://schemas.microsoft.com/ado/2007/08/
dataservices/related/customer"
type = "application/atom+xml;type=entry"
title = "customer" href = "Order(orderId=5001,
customer_customerId='IBM')/customer"/>
    <link rel = "http://schemas.microsoft.com/ado/2007/08/
dataservices/related/orderDetails"
type = "application/atom+xml;type=feed"
title = "orderDetails" href = "Order(orderId=5001,
customer_customerId='IBM')/orderDetails"/>
    <content type = "application/xml">
        <m:properties>
            <d:orderId m:type = "Edm.Int32">5001</d:orderId>
            <d:customer_customerId>IBM</d:customer_customerId>
            <d:orderDate m:type = "Edm.DateTime">2009-12-16T19:50:
11.125</d:orderDate>
            <d:shipCity>Rochester</d:shipCity>
            <d:shipCountry m:null = "true"/>
            <d:version m:type = "Edm.Int32">0</d:version>
        </m:properties>
    </content>
</entry>
</feed>

```

- Response Code: 200 OK

JSON

- Method: GET
- Request URI: http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/ Customer('IBM')/orders
- Request Header: Accept: application/json
- Request Payload: None
- Response Header: Content-Type: application/json
- Response Payload:

```

{"d": [{"__metadata": {"uri": "http://localhost:8080/wxsrestservice/
restservice/NorthwindGrid/Order(orderId=5000,
customer_customerId='IBM')",
"type": "NorthwindGridModel.Order"},
"orderId": 5000,
"customer_customerId": "IBM",
"orderDate": "\/Date(1260992789562)\/",
"shipCity": "Rochester",
"shipCountry": null,
"version": 0,
"customer": {"__deferred": {"uri": "http://localhost:8080/
wxsrestservice/restservice/NorthwindGrid/Order(orderId=

```

```

    5000,customer_customerId='IBM')/customer"}},
"orderDetails":{"__deferred":{"uri":"http://localhost:8080/
wxsrestservice/restservice/NorthwindGrid/Order(orderId=
5000,customer_customerId='IBM')/orderDetails"}},
{"__metadata":{"uri":"http://localhost:8080/wxsrestservice/
restservice/NorthwindGrid/
Order(orderId=5001,
customer_customerId='IBM')",
"type":"NorthwindGridModel.Order"},
"orderId":5001,
"customer_customerId":"IBM",
"orderDate":"\\/Date(1260993011125)\\/",
"shipCity":"Rochester",
"shipCountry":null,
"version":0,
"customer":{"__deferred":{"uri":"http://localhost:8080/
wxsrestservice/restservice/NorthwindGrid/Order(orderId=
5001,customer_customerId='IBM')/customer"}},
"orderDetails":{"__deferred":{"uri":"http://localhost:8080/
wxsrestservice/restservice/NorthwindGrid/Order(orderId=
5001,customer_customerId='IBM')/orderDetails"}}}]}}

```

- Response Code: 200 OK

Retrieve a property

A RetrievePrimitiveProperty request can be used to get the value of a property of an eXtreme Scale entity instance. The property value is represented as XML format for AtomPub requests and a JSON object for JSON requests in the response payload. For more details on RetrievePrimitiveProperty request, refer to MSDN: RetrievePrimitiveProperty Request.

The following RetrievePrimitiveProperty request example retrieves the contactName property of the Customer('IBM') entity.

AtomPub

- Method: GET
- Request URI: http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/ Customer('IBM')/contactName
- Request Header: Accept: application/xml
- Request Payload: None
- Response Header: Content-Type: application/atom+xml
- Response Payload:

```

<contactName xmlns:d="http://schemas.microsoft.com/ado/2007/08/dataservices">
  John Doe
</contactName>

```
- Response Code: 200 OK

JSON

- Method: GET
- Request URI: http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/ Customer('IBM')/contactName
- Request Header: Accept: application/json
- Request Payload: None
- Response Header: Content-Type: application/json
- Response Payload: {"d":{"contactName":"John Doe"}}
- Response Code: 200 OK

Retrieve a property value

A RetrieveValue request can be used to get the raw value of a property on an eXtreme Scale entity instance. The property value is represented as a raw value in the response payload. If the entity type is one of the following, then the media type of the response is "text/plain." Otherwise the response' media type is "application/octet-stream." These types are:

- Java primitive types and their respective wrappers
- java.lang.String
- byte[]
- Byte[]
- char[]
- Character[]
- enums
- java.math.BigInteger
- java.math.BigDecimal
- java.util.Date
- java.util.Calendar
- java.sql.Date
- java.sql.Time
- java.sql.Timestamp

For more details on the RetrieveValue request, refer to MSDN: RetrieveValue Request.

The following RetrieveValue request example retrieves the raw value of the contactName property of the Customer('IBM') entity.

- Request Method: GET
- Request URI: http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Customer('IBM')/contactName/\$value
- Request Header: Accept: text/plain
- Request Payload: None
- Response Header: Content-Type: text/plain
- Response Payload: John Doe
- Response Code: 200 OK

Retrieve a link

A RetrieveLink Request can be used to get the link(s) representing a to-one association or to-many association. For the to-one association, the link is from one eXtreme Scale Entity instance to another, and the link is represented in the response payload. For the to-many association, the links are from one eXtreme Scale Entity instance to all others in a specified eXtreme Scale entity collection, and the response is represented as a set of links in the response payload. For more details on RetrieveLink request, refer to MSDN: RetrieveLink Request.

Here is a RetrieveLink request example. In this example, we retrieve the association between entity Order(orderId=5000,customer_customerId='IBM') and its customer. The response shows the Customer entity URI.

AtomPub

- Method: GET
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Order(orderId=5000,customer_customerId='IBM')/$links/customer`
- Request Header: Accept: application/xml
- Request Payload: None
- Response Header: Content-Type: application/xml
- Response Payload:


```
<?xml version="1.0" encoding="utf-8"?>
<uri>http://localhost:8080/wxsrestservice/restservice/
  NorthwindGrid/Customer('IBM')</uri>
```
- Response Code: 200 OK

JSON

- Method: GET
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Order(orderId=5000,customer_customerId='IBM')/$links/customer`
- Request Header: Accept: application/json
- Request Payload: None
- Response Header: Content-Type: application/json
- Response Payload: `{"d":{"uri":"http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Customer('IBM')"}}`

Retrieve service metadata

A RetrieveServiceMetadata Request can be used to get the conceptual schema definition language (CSDL) document, which describes the data model associated with the eXtreme Scale REST data service. For more details on RetrieveServiceMetadata request, refer to MSDN: RetrieveServiceMetadata Request.

Retrieve service document

A RetrieveServiceDocument Request can be used to retrieve the Service Document describing the collection of resources exposed by the eXtreme Scale REST data service. For more details on RetrieveServiceDocument request, refer to MSDN: RetrieveServiceDocument Request.

Insert requests with REST data services

An InsertEntity Request can be used to insert a new eXtreme Scale entity instance, potentially with new related entities, into the eXtreme Scale REST data service.

Insert entity request

An InsertEntity Request can be used to insert a new eXtreme Scale entity instance, potentially with new related entities, into the eXtreme Scale REST data service. When inserting an entity, the client may specify if the resource or entity should be automatically linked to other existing entities in the data service.

The client must include the required binding information in the representation of the associated relation in the request payload.

In addition to supporting the insertion of a new EntityType instance (E1), the InsertEntity request also allows inserting new entities related to E1 (described by an entity relation) in a single Request. For example, when inserting a

Customer('IBM'), we can insert all the orders with Customer('IBM'). This form of an InsertEntity Request is also known as a *deep insert*. With a deep insert, the related entities must be represented using the inline representation of the relation associated with E1 that identifies the link to the to-be-inserted related entities.

The properties of the entity to be inserted are specified in the request payload. The properties are parsed by the REST data service and then set to the correspondent property on the entity instance. For the AtomPub format, the property is specified as a <d:PROPERTY_NAME> XML element. For JSON, the property is specified as a property of a JSON object.

If a property is missing in the request payload, then the REST data service sets the entity property value to the java default value. However, the database backend might reject such a default value, for example, if the column is not nullable in the database. Then a 500 response code will be returned to indicate an Internal Server error.

If there are duplicate properties specified in the payload, the last property will be used. All the previous values for the same property name are ignored by the REST data service.

If the payload contains a non-existent property, then the REST data service returns a 400 (Bad Request) response code to indicate the request sent by the client was syntactically incorrect.

If the key properties are missing, then the REST data service returns a response code of 400 (Bad Request) to indicate a missing key property.

If the payload contains a link to a related entity with a non-existent key, then the REST data service returns a 404 (Not Found) response code to indicate the linked entity cannot be found.

If the payload contains a link to a related entity with an incorrect association name, then the REST data service returns a 400 (Bad Request) response code to indicate the link cannot be found.

If the payload contains more than one link to a to-one relation, the last link will be used. All the previous links for the same association are ignored.

For more details on the InsertEntity request, see MSDN Library: InsertEntity Request.

An InsertEntity request inserts a Customer entity with key 'IBM'.

AtomPub

- Method: POST
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Customer('IBM')`
- Request Header: `Accept: application/atom+xml Content-Type: application/atom+xml`
- Request Payload:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<entry xmlns:d="http://schemas.microsoft.com/ado/2007/08/dataservices"
  xmlns:m="http://schemas.microsoft.com/ado/2007/08/dataservices/metadata"
  xmlns="http://www.w3.org/2005/Atom">
```

```

<category term="NorthwindGridModel.Customer"
  scheme="http://schemas.microsoft.com/ado/2007/08/dataservices/scheme" />
<content type="application/xml">
  <m:properties>
    <d:customerId>Rational</d:customerId>
    <d:city>Rochester</d:city>
    <d:companyName>Rational</d:companyName>
    <d:contactName>John Doe</d:contactName>
    <d:country>USA</d:country>
  </m:properties>
</content>
</entry>

```

- Response Header: Content-Type: application/atom+xml
- Response Payload:

```

<?xml version="1.0" encoding="ISO-8859-1"?>
<entry xmlns:d="http://schemas.microsoft.com/ado/2007/08/dataservices"
  xmlns:m="http://schemas.microsoft.com/ado/2007/08/dataservices/metadata"
  xmlns="http://www.w3.org/2005/Atom">
  <category term="NorthwindGridModel.Customer"
    scheme="http://schemas.microsoft.com/ado/2007/08/dataservices/scheme" />
  <content type="application/xml">
    <m:properties>
      <d:customerId>Rational</d:customerId>
      <d:city>Rochester</d:city>
      <d:companyName>Rational</d:companyName>
      <d:contactName>John Doe</d:contactName>
      <d:country>USA</d:country>
    </m:properties>
  </content>
</entry>

```

Response Header:

Content-Type: application/atom+xml

Response Payload:

```

<?xml version="1.0" encoding="utf-8"?>
<entry xml:base = "http://localhost:8080/wxsrestservice/restservice" xmlns:d =
  "http://schemas.microsoft.com/ado/2007/08/dataservices" xmlns:m =
  "http://schemas.microsoft.com/
  ado/2007/08/dataservices/metadata" xmlns = "http://www.w3.org/2005/Atom">
  <category term = "NorthwindGridModel.Customer" scheme = "http://schemas.
    microsoft.com/ado/2007/08/dataservices/scheme"/>
  <id>http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/
    Customer('Rational')</id>
  <title type = "text"/>
  <updated>2009-12-16T23:25:50.875Z</updated>
  <author>
    <name/>
  </author>
  <link rel = "edit" title = "Customer" href = "Customer('Rational')"/>
  <link rel = "http://schemas.microsoft.com/ado/2007/08/dataservices/related/
    orders" type = "application/atom+xml;type=feed"
    title = "orders" href = "Customer('Rational')/orders"/>
  <content type = "application/xml">
    <m:properties>
      <d:customerId>Rational</d:customerId>
      <d:city>Rochester</d:city>
      <d:companyName>Rational</d:companyName>
      <d:contactName>John Doe</d:contactName>
      <d:country>USA</d:country>
      <d:version m:type = "Edm.Int32">0</d:version>
    </m:properties>
  </content>
</entry>

```

- Response Code: 201 Created

JSON

- Method: POST
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/ Customer`
- Request Header: `Accept: application/json Content-Type: application/json`
- Request Payload:


```
{ "customerId": "Rational",
  "city": null,
  "companyName": "Rational",
  "contactName": "John Doe",
  "country": "USA", }
```
- Response Header: `Content-Type: application/json`
- Response Payload:


```
{ "d": { "__metadata": { "uri": "http://localhost:8080/wxsrestservice/restservice/
NorthwindGrid/Customer('Rational')",
"type": "NorthwindGridModel.Customer" },
  "customerId": "Rational",
  "city": null,
  "companyName": "Rational",
  "contactName": "John Doe",
  "country": "USA",
  "version": 0,
  "orders": { "__deferred": { "uri": "http://localhost:8080/wxsrestservice/restservice/
NorthwindGrid/Customer('Rational')/orders" } } }
```
- Response Code: 201 Created

Insert link request

An InsertLink Request can be used to create a new Link between two eXtreme Scale entity instances. The URI of the request must resolve to an eXtreme Scale to-many association. The payload of the request contains a single link which points to the to-many association target entity.

If the URI of the InsertLink request represents a to-one association, the REST data service returns a 400 (Bad request) response.

If the URI of the InsertLink request points to an association which does not exist, the REST data service returns a 404 (Not Found) response to indicate the link cannot be found.

If the payload contains a link with a key which does not exist, the REST data service returns a 404 (Not Found) response to indicate the linked entity cannot be found.

If the payload contains more than one link, the eXtreme Scale Rest Data Service will parse the first link. The remaining links are ignored.

For more details on InsertLink request, refer to: MSDN Library: InsertLink Request.

The following InsertLink request example creates a link from Customer('IBM') to Order(orderId=5000,customer_customerId='IBM').

AtomPub

- Method: POST
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/ Customer('IBM')/$link/orders`
- Request Header: `Content-Type: application/xml`

- Request Payload:


```
<?xml version="1.0" encoding="ISO-8859-1"?>
<uri>http://host:1000/wxsrestservice/restservice/NorthwindGrid/Order(orderId=
5000,customer_customerId='IBM')</uri>
```
- Response Payload: None
- Response Code: 204 No Content

JSON

- Method: POST
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Customer('IBM')/$links/orders`
- Request Header: Content-Type: application/json
- Request Payload:


```
{"uri": "http://host:1000/wxsrestservice/restservice/NorthwindGrid/Order(orderId
=5000,customer_customerId='IBM')"}

```
- Response Payload: None
- Response Code: 204 No Content

Update requests with REST data services

The WebSphere eXtreme Scale REST data service supports update requests for entities, entity primitive properties, and so on.

Update an entity

An UpdateEntity Request can be used to update an existing eXtreme Scale entity. The client can use an HTTP PUT method to replace an existing eXtreme Scale entity, or use an HTTP MERGE method to merge the changes into an existing eXtreme Scale entity.

When updating the entity, the client may specify if the entity, in addition to being updated, should be automatically linked to other existing entities in the data service that are related through single valued to-one associations.

The property of the entity to be updated is in the request payload. The property is parsed by the REST data service and then set to the correspondent property on the entity. For the AtomPub format, the property is specified as a `<d:PROPERTY_NAME>` XML element. For JSON, the property is specified as a property of a JSON object.

If a property is missing in the request payload, the REST data service sets the entity property value to the java default value for HTTP PUT method. However, the database backend might reject such a default value if, for example, the column is not nullable in the database. Then a 500 (Internal Server Error) response code will be returned to indicate an Internal Server Error. If a property is missing in the HTTP MERGE request payload, the REST data service will not change the existing property value.

If there are duplicate properties specified in the payload, the last property will be used. All the previous values with the same property name are ignored by the REST data service.

If the payload contains a non-existent property, the REST data service returns a 400 (Bad Request) response code to indicate the request sent by the client was syntactically incorrect.

As part of the serialization of a resource, if the payload of an Update request contains any of the key properties for the entity, the REST data service ignores those key values since entity keys are immutable.

For details on UpdateEntity request, refer to: MSDN Library: UpdateEntity Request.

An UpdateEntity request updates the city name of Customer('IBM') to 'Raleigh'.

AtomPub

- Method: PUT
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/ Customer('IBM')`
- Request Header: Content-Type: `application/atom+xml`
- Request Payload:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<entry xmlns:d="http://schemas.microsoft.com/ado/2007/08/dataservices"
xmlns:m="http://schemas.microsoft.com/ado/2007/08/dataservices/metadata"
xmlns="http://www.w3.org/2005/Atom">
<category term="NorthwindGridModel.Customer"
scheme="http://schemas.microsoft.com/ado/2007/08/dataservices/scheme" />
<title />
<updated>2009-07-28T21:17:50.609Z</updated>
<author>
<name />
</author>
<id />
<content type="application/xml">
<m:properties>
<d:customerId>IBM</d:customerId>
<d:city>Raleigh</d:city>
<d:companyName>IBM Corporation</d:companyName>
<d:contactName>Big Blue</d:contactName>
<d:country>USA</d:country>
</m:properties>
</content>
</entry>
```
- Response Payload: None
- Response Code: 204 No Content

JSON

- Method: PUT
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/ Customer('IBM')`
- Request Header: Content-Type: `application/json`
- Request Payload:

```
{"customerId": "IBM",
"city": "Raleigh",
"companyName": "IBM Corporation",
"contactName": "Big Blue",
"country": "USA",}
```
- Response Payload: None
- Response Code: 204 No Content

Update an entity primitive property

The UpdatePrimitiveProperty Request can update a property value of an eXtreme Scale entity. The property and value to be updated are in the request payload. The property cannot be a key property since eXtreme Scale does not allow clients to change entity keys.

For more details on the UpdatePrimitiveProperty request, refer to: MSDN Library: UpdatePrimitiveProperty Request.

Here is an UpdatePrimitiveProperty request example. In this example, we update the city name of Customer('IBM') to 'Raleigh'.

AtomPub

- Method: PUT
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Customer('IBM')/city`
- Request Header: Content-Type: application/xml
- Request Payload:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<city xmlns:d="http://schemas.microsoft.com/ado/2007/08/dataservices">
  Raleigh
</city>
```
- Response Payload: None
- Response Code: 204 No Content

JSON

- Method: PUT
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Customer('IBM')/city`
- Request Header: Content-Type: application/json
- Request Payload: `{"city":"Raleigh"}`
- Response Payload: None
- Response Code: 204 No Content

Update an entity primitive property value

The UpdateValue Request can update a raw property value of an eXtreme Scale entity. The value to be updated is represented as a raw value in the request payload. The property cannot be a key property since eXtreme Scale does not allow clients to change entity keys.

The content type of the request can be “text/plain” or “application/octet-stream” depending on the property type. Refer to section 6.3.1.4 for more details.

For more details on the UpdateValue request, refer to: MSDN Library: UpdateValue Request

Here is an UpdateValue request example. In this example we update the city name of Customer('IBM') to 'Raleigh'.

- Method: PUT
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Customer('IBM')/city/$value`

- Request Header: Content-Type: text/plain
- Request Payload: Raleigh
- Response Payload: None
- Response Code: 204 No Content

Update a link

The UpdateLink request can be used to establish an association between two eXtreme Scale entity instances. The association can be a single valued (to-one) relation or a multi valued (to-many) relation.

Updating a link between two eXtreme Scale entity instances can not only establish associations, but also remove associations. For example, if the client establishes a to-one association between an Order(orderId=5000,customer_customerId='IBM') entity and Customer('ALFKI') instance, it has to dissociate the Order(orderId=5000,customer_customerId='IBM') entity and entity from its currently associated Customer instance.

If either of the entity instances specified in the UpdateLink request cannot be found, the REST data service returns a 404 (Not Found) response.

If the URI of the UpdateLink request specifies a non-existent association, the REST data service returns a 404 (Not Found) response to indicate the link cannot be found.

If the URI specified in the UpdateLink request payload does not resolve to the same entity or the same key as specified in the URI, if exists, then the eXtreme Scale Rest Data Service returns a 400 (Bad Request) response.

If the UpdateLink request payload contains multiple links, then the REST data service will only parse the first link. The rest of the links are ignored.

For more details on the UpdateLink request, refer to: MSDN Library: UpdateLink Request.

Here is an UpdateLink request example. In this example, we update the customer relation of Order(orderId=5000,customer_customerId='IBM') entity and from Customer('IBM') to Customer('IBM').

Remember: The previous example is for illustration only. Because all associations are typically key-associations for a partitioned grid, the link cannot be changed.

AtomPub

- Method: PUT
- Request URI: http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Order(101)/\$links/customer
- Request Header: Content-Type: application/xml
- Request Payload:


```
<?xml version="1.0" encoding="ISO-8859-1"?>
<uri>
  http://host:1000/wxsrestservice/restservice/NorthwindGrid/Customer('IBM')
</uri>
```
- Response Payload: None
- Response Code: 204 No Content

JSON

- Method: PUT
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Order(orderId=5000,customer_customerId='IBM')/$links/customer`
- Request Header: Content-Type: `application/xml`
- Request Payload: `{"uri": "http://host:1000/wxsrestservice/restservice/NorthwindGrid/Customer('IBM')"}`
- Response Payload: None
- Response Code: 204 No Content

Delete requests with REST data services

The WebSphere eXtreme Scale REST data service can delete entities, property values and links.

Delete an entity

The DeleteEntity Request can delete an eXtreme Scale entity from the REST data service.

If any relation to the to-be-deleted entity has cascade-delete set, then the eXtreme Scale Rest data service will delete the related entity or entities. For more details on the DeleteEntity request, refer to MSDN Library: DeleteEntity Request.

The following DeleteEntity request deletes the customer with key 'IBM'.

- Method: DELETE
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/Customer('IBM')`
- Request Payload: None
- Response Payload: None
- Response Code: 204 No Content

Delete a property value

The DeleteValue Request sets an eXtreme Scale entity property to null.

Any property of an eXtreme Scale entity can be set to null with a DeleteValue request. To set a property to null, ensure all of the following:

- For any primitive number type and its wrapper, BigInteger, or BigDecimal, the property value is set to 0.
- For Boolean or boolean type, the property value is set to false.
- For char or Character type, the property value is set to character #X1 (NIL).
- For enum type, the property value is set to the enum value with ordinal 0.
- For all other types, the property value is set to null.

However, such a delete request could be rejected by the database backend if, for example, the property is not nullable in the database. In this case, the REST data service returns a 500 (Internal Server Error) response. For more details on the DeleteValue request, refer to: MSDN Library: DeleteValue Request.

Here is a DeleteValue request example. In this example, we set the contact name of Customer('IBM') to null.

- Method: DELETE
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/ Customer('IBM')/contactName`
- Request Payload: None
- Response Payload: None
- Response Code: 204 No Content

Delete a link

The DeleteLink request can remove an association between two eXtreme Scale entity instances. The association can be a to-one relation or a to-many relation. However, such a delete request could be rejected by the database backend if, for example, the foreign key constraint is set. In this case, the REST data service returns a 500 (Internal Server Error) response. For more details on the DeleteLink request, refer to: MSDN Library: DeleteLink Request.

The following DeleteLink request removes the association between Order(101) and its associated Customer.

- Method: DELETE
- Request URI: `http://localhost:8080/wxsrestservice/restservice/NorthwindGrid/ Order(101)/$links/customer`
- Request Payload: None
- Response Payload: None
- Response Code: 204 No Content

Optimistic concurrency

The eXtreme Scale REST data service uses an optimistic locking model by using native HTTP headers: If-Match, If-None-Match, and ETag. These headers are sent in request and response messages to relay an entity's version information from the server to client and client to server.

For more details on Optimistic Concurrency, refer to MSDN Library: Optimistic Concurrency (ADO.NET).

The eXtreme Scale REST data service enables optimistic concurrency for an entity if a version attribute is defined in the entity schema for that entity. A version property can be defined in the entity schema by a @Version annotation for Java classes or a <version/> attribute for entities defined using an entity descriptor XML file. The eXtreme Scale REST data service automatically propagates the value of the version property to the client in the ETag header for single entity responses using an m:etag attribute in the payload for multiple entity XML responses, and an etag attribute in the payload for multiple entity JSON responses.

For more details on defining an eXtreme Scale entity schema, see "Defining an entity schema" on page 61.

Chapter 4. Programming with system APIs and plug-ins

A plug-in is a component that provides a function to the pluggable components, which include ObjectGrid and BackingMap. To most effectively use eXtreme Scale as an in-memory data grid or database processing space, you should carefully determine how best you can maximize performance with available plug-ins.

Introduction to plug-ins

A WebSphere eXtreme Scale plug-in is a component that provides a certain type of function to the pluggable components that include ObjectGrid and BackingMap. WebSphere eXtreme Scale provides several plug points to allow applications and cache providers to integrate with various data stores, alternative client APIs and to improve overall performance of the cache. The product ships with several default, prebuilt plug-ins, but you can also build custom plug-ins with the application.

All plug-ins are concrete classes that implement one or more eXtreme Scale plug-in interfaces. These classes are then instantiated and invoked by the ObjectGrid at appropriate times. The ObjectGrid and BackingMaps each allow custom plug-ins to be registered.

ObjectGrid plug-ins

The following plug-ins are available for an ObjectGrid instance. If the plug-in is server side only, the plug-ins are removed on the client ObjectGrid and BackingMap instances. The ObjectGrid and BackingMap instances are only on the server.

- **TransactionCallback:** A TransactionCallback plug-in provides transaction life cycle events. If the TransactionCallback plug-in is the built-in JPATxCallback (com.ibm.websphere.objectgrid.jpa.JPATxCallback) class implementation, then the plug-in is server side only. However, the subclasses of the JPATxCallback class are not server side only.
- **ObjectGridEventListener:** An ObjectGridEventListener plug-in provides ObjectGrid life cycle events for the ObjectGrid, shards, and transactions.
- **SubjectSource, ObjectGridAuthorization, SubjectValidation:** eXtreme Scale provides several security endpoints to allow custom authentication mechanisms to be integrated with eXtreme Scale. (Server side only)
- **MapAuthorization:** (Server side only)

Common ObjectGrid plug-in requirements

The ObjectGrid instantiates and initializes plug-in instances using JavaBeans conventions. All of the previous plug-in implementations have the following requirements:

- The plug-in class must be a top-level public class.
- The plug-in class must provide a public, no-argument constructor.
- The plug-in class must be available in the class path for both servers and clients (as appropriate).
- Attributes must be set using the JavaBeans style property methods.

- Plug-ins, unless specifically noted, are registered before ObjectGrid initializes and cannot be changed after the ObjectGrid is initialized.

BackingMap plug-ins

The following plug-ins are available for a BackingMap:

- **Evictor:** An evictor plug-in is a default mechanism is provided for evicting cache entries and a plug-in for creating custom evictors.
- **ObjectTransformer:** An ObjectTransformer plug-in allows you to serialize, deserialize, and copy objects in the cache.
- **OptimisticCallback:** An OptimisticCallback plug-in allows you to customize versioning and comparison operations of cache objects when you are using the optimistic lock strategy.
- **MapEventListener:** A MapEventListener plug-in provides callback notifications and significant cache state changes that occur for a BackingMap.
- **Indexing:** Use the indexing feature, which is represented by the MapIndexplug-in plug-in, to build an index or several indexes on a BackingMap map to support non-key data access.
- **Loader:** A Loader plug-in on an ObjectGrid map acts as a memory cache for data that is typically kept in a persistent store on either the same system or some other system. (Server side only)

Plug-in life cycles

Most plug-ins have both initialize and destroy methods or equivalent methods, in addition to the methods for which they were designed to operate. These specialized methods of each plug-in are available to be invoked at designated functional points. Both initialize and destroy methods define the life cycle of plug-ins, which are controlled by their "owner" objects. An owner object is the object that actually uses the given plug-in. An owner can be a grid client, server, or a backing map.

When owner objects are initializing, they will invoke the initialize method of their owned plug-ins. During the destroy cycle of owner objects, the destroy method of plug-ins will be consequently invoked also. For details on the specifics of initialize and destroy methods, along with other methods capable with each plug-in, refer to the topics relevant to each plug-in.

As an example, consider a distributed environment. Both the client-side ObjectGrids and the server-side ObjectGrids can have their own plug-ins. The life cycle of a client-side ObjectGrid and, therefore, its plug-in instances are independent from all server-side ObjectGrid and plug-in instances.

In such a distributed topology, say you have an ObjectGrid named "myGrid" defined in the objectGrid.xml file and configured with a customized ObjectGridEventListener named myObjectGridEventListener. The objectGridDeployment.xml file defines the deployment policy for the myGrid ObjectGrid. Both the objectGrid.xml and objectGridDeployment.xml files are used to start container servers. During the startup of the container server, the server side myGrid ObjectGrid instance is initialized and the initialize method of the myObjectGridEventListener instance that is owned by the myObjectGrid instance is invoked. After the container server is started, your application can connect to the server-side myGrid ObjectGrid instance and obtain a client-side instance.

When obtaining the client-side myGrid ObjectGrid instance, the client-side myGrid instance goes through its own initialization cycle and invokes the initialize method of its own client-side myObjectGridEventListener instance. This client side myObjectGridEventListener instance is independent from the server-side myObjectGridEventListener instance. Its life cycle is controlled by its owner, which is the client-side myGrid ObjectGrid instance.

If your application disconnects or destroys the client-side myGrid ObjectGrid instance, the destroy method of the owned client-side myObjectGridEventListener instance are invoked automatically. However, this has no impact on server-side myObjectGridEventListener instance. The destroy method of the server-side myObjectGridEventListener instance is only be invoked during the destroy cycle of the server-side myGrid ObjectGrid instance when stopping a container server. That is, when stopping a container server, the contained ObjectGrid instances are destroyed and the destroy method of all their owned plug-ins is invoked.

Although the previous example applies specifically to the case of a client and a server instance of an ObjectGrid, the owner of a plug-in can also be a BackingMap and you must be careful to determine your configurations for plug-ins that you may write based on these life cycle considerations.

Plug-ins for evicting cache objects

WebSphere eXtreme Scale provides a default mechanism for evicting cache entries and a plug-in for creating custom evictors. An evictor controls the membership of entries in each BackingMap. The default evictor uses a time to live (TTL) eviction policy for each BackingMap. If you provide a pluggable evictor mechanism, it typically uses an eviction policy that is based on the number of entries instead of on time.

TimeToLive property

A default TTL evictor is created with every backing map. The default evictor removes entries based on a time to live concept. This behavior is defined by the ttlType attribute, which has the following types.

None

Specifies that entries never expire and therefore are never removed from the map.

Creation time

Specifies that entries are evicted depending on when they were created.

If you are using the CREATION_TIME ttlType, the evictor evicts an entry when its time from creation equals its TimeToLive attribute value, which is set in milliseconds in your application configuration. If you set the TimeToLive attribute value to 10 seconds, the entry is automatically evicted ten seconds after it was inserted.

It is important to take caution when setting this value for the CREATION_TIME ttlType. This evictor is best used when reasonably high amounts of additions to the cache exist that are only used for a set amount of time. With this strategy, anything that is created is removed after the set amount of time.

The CREATION_TIME ttlType is useful in scenarios such as refreshing stock quotes every 20 minutes or less. Suppose a Web application obtains

stock quotes, and getting the most recent quotes is not critical. In this case, the stock quotes are cached in an ObjectGrid for 20 minutes. After 20 minutes, the ObjectGrid map entries expire and are evicted. Every twenty minutes or so, the ObjectGrid map uses the Loader plug-in to refresh the map data with fresh data from the database. The database is updated every 20 minutes with the most recent stock quotes.

Last access time

Specifies that entries are evicted depending upon when they were last accessed, whether they were read or updated.

Last update time

Specifies that entries are evicted depending upon when they were last updated.

If you are using the `LAST_ACCESS_TIME` or `LAST_UPDATE_TIME` ttlType attribute, set the `TimeToLive` to a lower number than if you are using the `CREATION_TIME` ttlType, because the entries `TimeToLive` attribute is reset every time it is accessed. In other words, if the `TimeToLive` attribute is equal to 15 and an entry has existed for 14 seconds but then gets accessed, it does not expire again for another 15 seconds. If you set the `TimeToLive` to a relatively high number, many entries might never be evicted. However, if you set the value to something like 15 seconds, entries might be removed when they are not often accessed.

The `LAST_ACCESS_TIME` or `LAST_UPDATE_TIME` ttlType is useful in scenarios such as holding session data from a client, using an ObjectGrid map. Session data must be destroyed if the client does not use the session data for some period of time. For example, the session data times out after 30 minutes of no activity by the client. In this case, using a TTL type of `LAST_ACCESS_TIME` or `LAST_UPDATE_TIME` with the `TimeToLive` attribute set to 30 minutes is appropriate for this application.

The following example creates a backing map, set its default evictor ttlType attribute, and sets its `TimeToLive` property.

```
ObjectGrid objGrid = new ObjectGrid();
BackingMap bMap = objGrid.defineMap("SomeMap");
bMap.setTtlEvictorType(TTLType.LAST_ACCESS_TIME);
bMap.setTimeToLive(1800);
```

Most evictor settings should be set before you initialize the ObjectGrid.

You may also write your own evictors: For more information, see the information about writing a custom evictor in the *Programming Guide*.

Optional evictors

The default TTL evictor uses an eviction policy that is based on time, and the number of entries in the BackingMap has no affect on the expiration time of an entry. You can use an optional pluggable evictor to evict entries based on the number of entries that exist instead of based on time.

The following optional pluggable evictors provide some commonly used algorithms for deciding which entries to evict when a BackingMap grows beyond some size limit:

- The `LRUEvictor` evictor uses a least recently used (LRU) algorithm to decide which entries to evict when the BackingMap exceeds a maximum number of entries.

- The LFUevictor evictor uses a least frequently used (LFU) algorithm to decide which entries to evict when the BackingMap exceeds a maximum number of entries.

The BackingMap informs an evictor as entries are created, modified, or removed in a transaction. The BackingMap keeps track of these entries and chooses when to evict one or more entries from the BackingMap.

A BackingMap has no configuration information for a maximum size. Instead, evictor properties are set to control the evictor behavior. Both the LRUEvictor and the LFUevictor have a maximum size property that is used to cause the evictor to begin to evict entries after the maximum size is exceeded. Like the TTL evictor, the LRU and LFU evictors might not immediately evict an entry when the maximum number of entries is reached to minimize impact on performance.

If the LRU or LFU eviction algorithm is not adequate for a particular application, you can write your own evictors to create your eviction strategy.

Memory-based eviction

Important: Memory-based eviction is only supported on Java Platform, Enterprise Edition Version 5 or later.

All built-in evictors support memory-based eviction that can be enabled on the BackingMap interface by setting the evictionTriggers attribute of BackingMap to MEMORY_USAGE_THRESHOLD. For more information about how to set the evictionTriggers attribute on BackingMap, see the information about the BackingMap interface and the ObjectGrid descriptor XML file in the *Programming Guide*.

Memory-based eviction is based on heap usage threshold. When memory-based eviction is enabled on BackingMap and the BackingMap has any built-in evictor, the usage threshold is set to a default percentage of total memory if the threshold has not been previously set.

When you are using memory-based eviction, you should configure the garbage collection threshold to the same value as their target heap utilization. For example, if the memory-based eviction threshold is set at 50 percent and the garbage collection threshold is at the default 70 percent level, then the heap utilization can go as high as 70 percent. This heap utilization increase occurs because memory-based eviction is only triggered after a garbage collection cycle.

The memory-based eviction algorithm used by WebSphere eXtreme Scale is sensitive to the behavior of the garbage collection algorithm in use. The best algorithm for memory-based eviction is the IBM default throughput collector. Generation garbage collection algorithms can cause undesired behavior, and so you should not use these algorithms with memory-based eviction.

To change the usage threshold percentage, set the memoryThresholdPercentage property on the container and server property files for eXtreme Scale server processes.

During runtime, if the memory usage exceeds the target usage threshold, memory-based evictors start evicting entries and try to keep memory usage below the target usage threshold. However, no guarantee exists that the eviction speed is

fast enough to avoid a potential out of memory error if the system runtime continues to quickly consume memory.

TimeToLive (TTL) evictor

WebSphere eXtreme Scale provides a default mechanism for evicting cache entries and a plug-in for creating custom evictors. An evictor controls the membership of entries in each BackingMap instance.

Enable the TTL evictor programmatically

TTL evictors are associated with BackingMap instances. The default evictor uses a time-to-live (TTL) eviction policy for each BackingMap instance. If you provide a pluggable evictor mechanism, it typically uses an eviction policy that is based on the number of entries instead of on time.

The following snippet of code uses the BackingMap interface to set the expiration time for each entry to 10 minutes after the entry was created.

programmatic time-to-live evictor

```
import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
import com.ibm.websphere.objectgrid.ObjectGridManager;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.BackingMap;
import com.ibm.websphere.objectgrid.TTLType;
ObjectGridManager ogManager = ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid og = ogManager.createObjectGrid( "grid" );
BackingMap bm = og.defineMap( "myMap" );
bm.setTtlEvictorType( TTLType.CREATION_TIME );
bm.setTimeToLive( 600 );
```

The setTimeToLive method argument is 600 because it indicates the time-to-live value is in seconds. The preceding code must run before the initialize method is invoked on the ObjectGrid instance. These BackingMap attributes cannot be changed after the ObjectGrid instance is initialized. After the code runs, any entry that is inserted into the myMap BackingMap has an expiration time. After the expiration time is reached, the TTL evictor removes the entry.

To set the expiration time to the last access time plus 10 minutes, change the argument that is passed to the setTtlEvictorType method from TTLType.CREATION_TIME to TTLType.LAST_ACCESS_TIME. With this value, the expiration time is computed as the last access time plus 10 minutes. When an entry is first created, the last access time is the creation time. To base the expiration time on the last *update*, instead of merely the last *access* (whether or not it involved an update), substitute the TTLType.LAST_UPDATE_TIME setting for the TTLType.LAST_ACCESS_TIME setting.

When using the TTLType.LAST_ACCESS_TIME or TTLType.LAST_UPDATE_TIME setting, you can use the ObjectMap and JavaMap interfaces to override the BackingMap time-to-live value. This mechanism allows an application to use a different time-to-live value for each entry that is created. Assume that the preceding snippet of code set the ttlType attribute to LAST_ACCESS_TIME and set the time-to-live value to 10 minutes. An application can then override the time-to-live value for each entry by running the following code prior to creating or modifying an entry:

```
import com.ibm.websphere.objectgrid.Session;
import com.ibm.websphere.objectgrid.ObjectMap;
Session session = og.getSession();
ObjectMap om = session.getMap( "myMap" );
```



```
int oldTimeToLive1 = om.setTimeToLive( 1800 );
om.insert("key1", "value1" );
int oldTimeToLive2 = om.setTimeToLive( 1200 );
om.insert("key2", "value2" );
```

In the previous snippet of code, the entry with the key1 key has an expiration time of the insert time plus 30 minutes as a result of the setTimeToLive(1800) method invocation on the ObjectMap instance. The oldTimeToLive1 variable is set to 600 because the time-to-live value from the BackingMap is used as a default value if the setTimeToLive method was not previously called on the ObjectMap instance.

The entry with the key2 key has an expiration time of insert time plus 20 minutes as a result of the setTimeToLive(1200) method call on the ObjectMap instance. The oldTimeToLive2 variable is set to 1800 because the time-to-live value from the previous ObjectMap.setTimeToLive method invocation set the time-to-live value to 1800.

The previous example shows two map entries being inserted in the myMap map for keys key1 and key2. At a later point in time, the application can still update these map entries while retaining the time-to-live values that are used at insert time for each map entry. The following example illustrates how to retain the time-to-live values by using a constant defined in the ObjectMap interface:

```
Session session = og.getSession();
ObjectMap om = session.getMap( "myMap" );
om.setTimeToLive( ObjectMap.USE_DEFAULT );
session.begin();
om.update("key1", "updated value1" );
om.update("key2", "updated value2" );
om.insert("key3", "value3" );
session.commit();
```

Since the ObjectMap.USE_DEFAULT special value is used on the setTimeToLive method call, the key1 key retains its time-to-live value of 1800 seconds and the key2 key retains its time-to-live value of 1200 seconds because those values were used when these map entries were inserted by the prior transaction.

The previous example also shows a new map entry for the key3 key insert. In this case, the USE_DEFAULT special value indicates to use the default setting of time-to-live value for this map. The default value is defined by the time-to-live BackingMap attribute. See BackingMap interface attributes for information about how the time-to-live attribute is defined on the BackingMap instance.

See the API documentation for the setTimeToLive method on the ObjectMap and JavaMap interfaces. The documentation explains that an IllegalStateException exception results if the BackingMap.getTtlEvictorType method returns anything other than the TTLType.LAST_ACCESS_TIME or TTLType.LAST_UPDATE_TIME value. The ObjectMap and JavaMap interfaces can override the time-to-live value only when you are using the LAST_ACCESS_TIME or TTLType.LAST_UPDATE_TIME setting for the TTL evictor type. The setTimeToLive method cannot be used to override the time-to-live value when you are using the evictor type setting CREATION_TIME or NONE.

Enable the TTL evictor using XML configuration

Instead of using the BackingMap interface to programmatically set the BackingMap attributes to be used by the TTL evictor, you can use an XML file to configure each BackingMap instance. The following code demonstrates how to set these attributes for three different BackingMap maps:

enabling time-to-live evictor using XML

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">
<objectGrids>
  <objectGrid name="grid1">
    <backingMap name="map1" ttlEvictorType="NONE" />
    <backingMap name="map2" ttlEvictorType="LAST_ACCESS_TIME|LAST_UPDATE_TIME"
      timeToLive="1800" />
    <backingMap name="map3" ttlEvictorType="CREATION_TIME"
      timeToLive="1200" />
  </objectGrid>
</objectGrids>
```

The preceding example shows that the map1 BackingMap instance uses a NONE TTL evictor type. The map2 BackingMap instance uses either a LAST_ACCESS_TIME or LAST_UPDATE_TIME TTL evictor type – specify only one or the other of these settings – and has a time-to-live value of 1800 seconds, or 30 minutes. The map3 BackingMap instance is defined to use a CREATION_TIME TTL evictor type and has a time-to-live value of 1200 seconds, or 20 minutes.

Plug in a pluggable evictor

Since evictors are associated with BackingMaps, use the BackingMap interface to specify the pluggable evictor.

Optional pluggable evictors

The default TTL evictor uses an eviction policy that is based on time, and the number of entries in the BackingMap has no affect on the expiration time of an entry. You can use an optional pluggable evictor to evict entries based on the number of entries that exist instead of based on time.

The following optional pluggable evictors provide some commonly used algorithms for deciding which entries to evict when a BackingMap grows beyond some size limit.

- The LRUEvictor evictor uses a least recently used (LRU) algorithm to decide which entries to evict when the BackingMap exceeds a maximum number of entries.
- The LFUEvictor evictor uses a least frequently used (LFU) algorithm to decide which entries to evict when the BackingMap exceeds a maximum number of entries.

The BackingMap informs an evictor as entries are created, modified, or removed in a transaction. The BackingMap keeps track of these entries and chooses when to evict one or more entries from the BackingMap instance.

A BackingMap instance has no configuration information for a maximum size. Instead, evictor properties are set to control the evictor behavior. Both the LRUEvictor and the LFUEvictor have a maximum size property that is used to cause the evictor to begin to evict entries after the maximum size is exceeded. Like the TTL evictor, the LRU and LFU evictors might not immediately evict an entry when the maximum number of entries is reached to minimize impact on performance.

If the LRU or LFU eviction algorithm is not adequate for a particular application, you can write your own evictors to create your eviction strategy.

Using optional pluggable evictors

To add optional pluggable evictors into the BackingMap configuration, you can use programmatic configuration or XML configuration, as described in the following section.

Programmatically plug in a pluggable evictor

Because evictors are associated with BackingMaps, use the BackingMap interface to specify the pluggable evictor. The following code snippet is an example of specifying a LRUEvictor evictor for the map1 BackingMap and a LFUEvictor evictor for the map2 BackingMap instance:

plugging in an evictor programmatically

```
import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
import com.ibm.websphere.objectgrid.ObjectGridManager;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.BackingMap;
import com.ibm.websphere.objectgrid.plugins.builtins.LRUEvictor;
import com.ibm.websphere.objectgrid.plugins.builtins.LFUEvictor;
ObjectGridManager ogManager = ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid og = ogManager.createObjectGrid( "grid" );
BackingMap bm = og.defineMap( "map1" );
LRUEvictor evictor = new LRUEvictor();
evictor.setMaxSize(1000);
evictor.setSleepTime( 15 );
evictor.setNumberOfLRUQueues( 53 );
bm.setEvictor(evictor);
bm = og.defineMap( "map2" );
LFUEvictor evictor2 = new LFUEvictor();
evictor2.setMaxSize(2000);
evictor2.setSleepTime( 15 );
evictor2.setNumberOfHeaps( 211 );
bm.setEvictor(evictor2);
```

The preceding snippet shows an LRUEvictor evictor being used for map1 BackingMap with an approximate maximum number of entries of 53,000 (53 * 1000). The LFUEvictor evictor is used for the map2 BackingMap with an approximate maximum number of entries of 422,000 (211 * 2000). Both the LRU and LFU evictors have a sleep time property that indicates how long the evictor sleeps before waking up and checking to see if any entries need to be evicted. The sleep time is specified in seconds. A value of 15 seconds is a good compromise between performance impact and preventing BackingMap from growing too large. The goal is to use the largest sleep time possible without causing the BackingMap to grow to an excessive size.

The setNumberOfLRUQueues method sets the LRUEvictor property that indicates how many LRU queues the evictor uses to manage LRU information. A collection of queues is used so that every entry does not keep LRU information in the same queue. This approach can improve performance by minimizing the number of map entries that need to synchronize on the same queue object. Increasing the number of queues is a good way to minimize the impact that the LRU evictor can cause on performance. A good starting point is to use ten percent of the maximum number of entries as the number of queues. Using a prime number is typically better than using a number that is not prime. The setMaxSize method indicates how many entries are allowed in each queue. When a queue reaches its maximum number of

entries, the least recently used entry or entries in that queue are evicted the next time that the evictor checks to see if any entries need to be evicted.

The `setNumberOfHeaps` method sets the `LFUEvictor` property to set how many binary heap objects the `LFUEvictor` uses to manage LFU information. Again, a collection is used to improve performance. Using ten percent of the maximum number of entries is a good starting point and a prime number is typically better than using a number that is not prime. The `setMaxSize` method indicates how many entries are allowed in each heap. When a heap reaches its maximum number of entries, the least frequently used entry or entries in that heap are evicted the next time that the evictor checks to see if any entries need to be evicted.

XML configuration approach to plug in a pluggable evictor

Instead of using various APIs to programmatically plug in an evictor and set its properties, an XML file can be used to configure each `BackingMap` as illustrated in the following sample:

```
plugging in an evictor using XML
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
xmlns="http://ibm.com/ws/objectgrid/config">
<objectGrids>
  <objectGrid name="grid">
    <backingMap name="map1" ttlEvictorType="NONE" pluginCollectionRef="LRU" />
    <backingMap name="map2" ttlEvictorType="NONE" pluginCollectionRef="LFU" />
  </objectGrid>
</objectGrids>
<backingMapPluginCollections>
  <backingMapPluginCollection id="LRU">
    <bean id="Evictor" className="com.ibm.websphere.objectgrid.plugins.builtins.LRUEvictor">
      <property name="maxSize" type="int" value="1000" description="set max size
for each LRU queue" />
      <property name="sleepTime" type="int" value="15" description="evictor
thread sleep time" />
      <property name="numberOfLRUQueues" type="int" value="53" description="set number
of LRU queues" />
    </bean>
  </backingMapPluginCollection>
  <backingMapPluginCollection id="LFU">
    <bean id="Evictor" className="com.ibm.websphere.objectgrid.plugins.builtins.LFUEvictor">
      <property name="maxSize" type="int" value="2000" description="set max size for each LFU heap" />
      <property name="sleepTime" type="int" value="15" description="evictor thread sleep time" />
      <property name="numberOfHeaps" type="int" value="211" description="set number of LFU heaps" />
    </bean>
  </backingMapPluginCollection>
</backingMapPluginCollections>
</objectGridConfig>
```

Memory-based eviction

All built-in evictors support memory-based eviction that can be enabled on `BackingMap` interface by setting the `evictionTriggers` attribute of `BackingMap` to `"MEMORY_USAGE_THRESHOLD"`. For more information about how to set the `evictionTriggers` attribute on `BackingMap`, see `BackingMap` interface and `eXtreme Scale` configuration reference.

Memory-based eviction is based on heap usage threshold. When memory-based eviction is enabled on `BackingMap` and the `BackingMap` has any built-in evictor, the usage threshold is set to a default percentage of total memory if the threshold has not been previously set.

To change the default usage threshold percentage, set the `memoryThresholdPercentage` property on container and server property file for `eXtreme Scale` server process. To set the target usage threshold on an `eXtreme Scale` client process, you can use the `MemoryPoolMXBean`. See also: `containerServer.props` file and Starting `eXtreme Scale` server processes.

During run time, if the memory usage exceeds the target usage threshold, memory-based evictors start evicting entries and try to keep memory usage below the target usage threshold. However, there is no guarantee that the eviction speed is fast enough to avoid a potential out of memory error if the system run time continues to quickly consume memory.

Writing a custom evictor

WebSphere eXtreme Scale allows you to write a custom eviction implementation.

You must create a custom evictor that implements the evictor interface and follows the common eXtreme Scale plug-in conventions. The interface follows:

```
public interface Evictor
{
    void initialize(BackingMap map, EvictionEventCallback callback);
    void activate();
    void apply(LogSequence sequence);
    void deactivate();
    void destroy();
}
```

- The initialize method is invoked during initialization of the BackingMap object. This method initializes an Evictor plug-in with a reference to the BackingMap and a reference to an object that implements the `com.ibm.websphere.objectgrid.plugins.EvictionEventCallback` interface.
- The activate method is called to activate the Evictor. After this method is called, the Evictor can use the EvictionEventCallback interface to evict map entries. If the Evictor attempts to use the EvictionEventCallback interface to evict map entries before the activate method is called, an `IllegalStateException` exception results.
- The apply method is invoked when transactions that access one or more entries of the BackingMap are committed. The apply method is passed a reference to an object that implements the `com.ibm.websphere.objectgrid.plugins.LogSequence` interface. The LogSequence interface allows an Evictor plug-in to determine which BackingMap entries were created, modified, or removed by the transaction. An Evictor uses this information in deciding when and which entries to evict.
- The deactivate method is called to deactivate the Evictor. After this method is called, the Evictor must stop using the EvictionEventCallback interface to evict map entries. If the Evictor uses the EvictionEventcallback interface after this method is called, an `IllegalStateException` exception results.
- The destroy method is invoked when the BackingMap is being destroyed. This method allows an Evictor to terminate any threads that it might have created.

The EvictionEventCallback interface has the following methods:

```
public interface EvictionEventCallback
{
    void evictMapEntries(List evictorDataList) throws ObjectGridException;
    void evictEntries(List keysToEvictList) throws ObjectGridException;
    void setEvictorData(Object key, Object data);
    Object getEvictorData(Object key);
}
```

The EvictionEventCallback methods are used by an Evictor plug-in to call back to the eXtreme Scale framework as follows:

- The setEvictorData method is used by an evictor to request the framework that is used to store and associate some evictor object it creates with the entry indicated by the key argument. The data is evictor specific and is determined by

the information the evictor needs to implement the algorithm it is using. For example, in a least frequently used algorithm, the evictor maintains a count in the evictor data object for tracking how many times the apply method is invoked with a LogElement that refers to an entry for a given key.

- The `getEvictorData` method is used by an evictor to retrieve the data it passed to the `setEvictorData` method during a prior apply method invocation. If evictor data for the specified key argument is not found, a special `KEY_NOT_FOUND` object that is defined on the `EvictorCallback` interface is returned.
- The `evictMapEntries` method is used by an evictor to request the eviction of one or more map entries. Each object in the `evictorDataList` parameter must implement the `com.ibm.websphere.objectgrid.plugins.EvictorData` interface. Also, the same `EvictorData` instance that is passed to the `setEvictorData` method must be in the evictor data list parameter of this method. The `getKey` method of the `EvictorData` interface is used to determine which map entry to evict. The map entry is evicted if the cache entry currently contains the exact same `EvictorData` instance that is in the evictor data list for this cache entry.
- The `evictEntries` method is used by an evictor to request eviction of one or more map entries. This method is used only if the object that is passed to the `setEvictorData` method does not implement the `com.ibm.websphere.objectgrid.plugins.EvictorData` interface.

After a transaction completes eXtreme Scale calls the apply method of the Evictor interface. All transaction locks that were acquired by the completed transaction are no longer held. Potentially, multiple threads can call the apply method at the same time, and each thread can complete its own transaction. Because transaction locks are already released by the completed transaction, the apply method must provide its own synchronization to ensure the apply method is thread safe.

The reason to implement the `EvictorData` interface and use the `evictMapEntries` method instead of the `evictEntries` method is to close a potential timing window. Consider the following sequence of events:

1. Transaction 1 completes and calls the apply method with a `LogSequence` that deletes the map entry for key 1.
2. Transaction 2 completes and calls the apply method with a `LogSequence` that inserts a new map entry for key 1. In other words, transaction 2 recreates the map entry that was deleted by transaction 1.

Because the evictor runs asynchronously from threads that run transactions, it is possible that when the evictor decides to evict key 1, it might be evicting either the map entry that existed prior to transaction 1 completion, or it might be evicting the map entry that was recreated by transaction 2. To eliminate timing windows and to eliminate uncertainty as to which version of the key 1 map entry the evictor intended to evict, implement the `EvictorData` interface by the object that is passed to the `setEvictorData` method. Use the same `EvictorData` instance for the life of a map entry. When that map entry is deleted and is then recreated by another transaction, the evictor should use a new instance of the `EvictorData` implementation. By using the `EvictorData` implementation and by using the `evictMapEntries` method, the evictor can ensure that the map entry is evicted if and only if the cache entry that is associated with the map entry contains the correct `EvictorData` instance.

The `Evictor` and `EvictorEventCallback` interfaces allow an application to plug in an evictor that implements a user-defined algorithm for eviction. The following snippet of code illustrates how you can implement the initialize method of `Evictor` interface:

```

import com.ibm.websphere.objectgrid.BackingMap;
import com.ibm.websphere.objectgrid.plugins.EvictionEventCallback;
import com.ibm.websphere.objectgrid.plugins.Evictor;
import com.ibm.websphere.objectgrid.plugins.LogElement;
import com.ibm.websphere.objectgrid.plugins.LogSequence;
import java.util.LinkedList;
// Instance variables
private BackingMap bm;
private EvictionEventCallback evictorCallback;
private LinkedList queue;
private Thread evictorThread;
public void initialize(BackingMap map, EvictionEventCallback callback)
{
    bm = map;
    evictorCallback = callback;
    queue = new LinkedList();
    // spawn evictor thread
    evictorThread = new Thread( this );
    String threadName = "MyEvictorForMap-" + bm.getName();
    evictorThread.setName( threadName );
    evictorThread.start();
}

```

The preceding code saves the references to the map and callback objects in instance variables so that they are available to the apply and destroy methods. In this example, a linked list is created that is used as a first in, first out queue for implementing a least recently used (LRU) algorithm. A thread is spawned off and a reference to the thread is kept as an instance variable. By keeping this reference, the destroy method can interrupt and terminate the spawned thread.

Ignoring synchronization requirements to make code thread safe, the following snippet of code illustrates how the apply method of the Evictor interface can be implemented:

```

import com.ibm.websphere.objectgrid.BackingMap;
import com.ibm.websphere.objectgrid.plugins.EvictionEventCallback;
import com.ibm.websphere.objectgrid.plugins.Evictor;
import com.ibm.websphere.objectgrid.plugins.EvictorData;
import com.ibm.websphere.objectgrid.plugins.LogElement;
import com.ibm.websphere.objectgrid.plugins.LogSequence;

public void apply(LogSequence sequence)
{
    Iterator iter = sequence.getAllChanges();
    while ( iter.hasNext() )
    {
        LogElement elem = (LogElement)iter.next();
        Object key = elem.getKey();
        LogElement.Type type = elem.getType();
        if ( type == LogElement.INSERT )
        {
            // do insert processing here by adding to front of LRU queue.
            EvictorData data = new EvictorData(key);
            evictorCallback.setEvictorData(key, data);
            queue.addFirst( data );
        }
        else if ( type == LogElement.UPDATE || type == LogElement.FETCH || type == LogElement.TOUCH )
        {
            // do update processing here by moving EvictorData object to
            // front of queue.
            EvictorData data = evictorCallback.getEvictorData(key);
            queue.remove(data);
            queue.addFirst(data);
        }
        else if ( type == LogElement.DELETE || type == LogElement.EVICT )
        {
            // do remove processing here by removing EvictorData object
            // from queue.
            EvictorData data = evictorCallback.getEvictorData(key);
            if ( data == EvictionEventCallback.KEY_NOT_FOUND )
            {
                // Assumption here is your asynchronous evictor thread
                // evicted the map entry before this thread had a chance
                // to process the LogElement request. So you probably
                // need to do nothing when this occurs.
            }
            else
            {
                // Key was found. So process the evictor data.
                if ( data != null )
                {
                    // Ignore null returned by remove method since spawned
                    // evictor thread may have already removed it from queue.
                    // But we need this code in case it was not the evictor
                    // thread that caused this LogElement to occur.
                }
            }
        }
    }
}

```

```

        queue.remove( data );
    }
    else
    {
        // Depending on how you write you Evictor, this possibility
        // may not exist or it may indicate a defect in your evictor
        // due to improper thread synchronization logic.
    }
}
}
}
}

```

Insert processing in the apply method typically handles the creation of an evictor data object that is passed to the setEvictorData method of the EvictionEventCallback interface. Because this evictor illustrates a LRU implementation, the EvictorData is also added to the front of the queue that was created by the initialize method. Update processing in the apply method typically updates the evictor data object that was created by some prior invocation of the apply method (for example, by the insert processing of the apply method). Because this evictor is an LRU implementation, it needs to move the EvictorData object from its current queue position to the front of the queue. The spawned evictor thread removes the last EvictorData object in the queue because the last queue element represents the least recently used entry. The assumption is that the EvictorData object has a getKey method on it so that the evictor thread knows the keys of the entries that need to be evicted. Keep in mind that this example is ignoring synchronization requirements to make code thread safe. A real custom evictor is more complicated because it deals with synchronization and performance bottlenecks that occur as a result of the synchronization points.

The following snippets of code illustrate the destroy method and the run method of the runnable thread that the initialize method spawned:

```

// Destroy method simply interrupts the thread spawned by the initialize method.
public void destroy()
{
    evictorThread.interrupt();
}

// Here is the run method of the thread that was spawned by the initialize method.
public void run()
{
    // Loop until destroy method interrupts this thread.
    boolean continueToRun = true;
    while ( continueToRun )
    {
        try
        {
            // Sleep for a while before sweeping over queue.
            // The sleepTime is a good candidate for a evictor
            // property to be set.
            Thread.sleep( sleepTime );
            int queueSize = queue.size();
            // Evict entries if queue size has grown beyond the
            // maximum size. Obviously, maximum size would
            // be another evictor property.
            int numToEvict = queueSize - maxSize;
            if ( numToEvict > 0 )
            {
                // Remove from tail of queue since the tail is the
                // least recently used entry.
                List evictList = new ArrayList( numToEvict );
                while( queueSize > ivMaxSize )
                {
                    EvictorData data = null;
                    try
                    {
                        EvictorData data = (EvictorData) queue.removeLast();
                        evictList.add( data );
                        queueSize = queue.size();
                    }
                    catch ( NoSuchElementException nse )
                    {
                        // The queue is empty.
                        queueSize = 0;
                    }
                }
            }
            // Request eviction if key list is not empty.
            if ( ! evictList.isEmpty() )

```



```

        {
            evictorCallback.evictMapEntries( evictList );
        }
    }
}
catch ( InterruptedException e )
{
    continueToRun = false;
}
} // end while loop
} // end run method.

```

Optional RollBackEvictor interface

The `com.ibm.websphere.objectgrid.plugins.RollbackEvictor` interface can be optionally implemented by an Evictor plug-in. By implementing this interface, an evictor can be invoked not only when transactions are committed, but also when transactions are rolled back.

```

public interface RollbackEvictor
{
    void rollingBack( LogSequence ls );
}

```

The `rollingBack` method is called only if a transaction is committed. If a transaction is rolled back and the `RollbackEvictor` interface is implemented by the evictor, the `rollingBack` method is invoked. If the `RollbackEvictor` interface is not implemented and the transaction rolls back, the `apply` method and the `rollingBack` method are not called.

Plug-in evictor performance best practices

If you use plug-in evictors, they are not active until you create them and associate them with a backing map. The following best practices will increase performance for least frequently used (LFU) and least recently used (LRU) evictors.

Least frequently used (LFU) evictor

The concept of a LFU evictor is to remove entries from the map that are used infrequently. The entries of the map are spread over a set amount of binary heaps. As the usage of a particular cache entry grows, it becomes ordered higher in the heap. When the evictor attempts a set of evictions it removes only the cache entries that are located lower than a specific point on the binary heap. As a result, the least frequently used entries are evicted.

Least recently used (LRU) evictor

The LRU Evictor follows the same concepts of the LFU Evictor with a few differences. The main difference is that the LRU uses a first in, first out queue (FIFO) instead of a set of binary heaps. Every time a cache entry is accessed, it moves to the head of the queue. Consequently, the front of the queue contains the most recently used map entries and the end becomes the least recently used map entries. For example, the A cache entry is used 50 times, and the B cache entry is used only once right after the A cache entry. In this situation, the B cache entry is at the front of the queue because it was used most recently, and the A cache entry is at the end of the queue. The LRU evictor evicts the cache entries that are at the tail of the queue, which are the least recently used map entries.

LFU and LRU properties and best practices to improve performance

Number of heaps

When using the LFU evictor, all of the cache entries for a particular map are ordered over the number of heaps that you specify, improving performance drastically and preventing all of the evictions from synchronizing on one binary heap that contains all of the ordering for the map. More heaps also speeds up the time that is required for reordering the heaps because each heap has fewer entries. Set the number of heaps to 10% of the number of entries in your BaseMap.

Number of queues

When using the LRU evictor, all of the cache entries for a particular map are ordered over the number of LRU queues that you specify, improving performance drastically and preventing all of the evictions from synchronizing on one queue that contains all of the ordering for the map. Set the number of queues to 10% of the number of entries in your BaseMap.

MaxSize property

When an LFU or LRU evictor begins evicting entries, it uses the MaxSize evictor property to determine how many binary heaps or LRU queue elements to evict. For example, assume that you set the number of heaps or queues to have about 10 map entries in each map queue. If your MaxSize property is set to 7, the evictor evicts 3 entries from each heap or queue object to bring the size of each heap or queue back down to 7. The evictor only evicts map entries from a heap or queue when that heap or queue has more than the MaxSize property value of elements in it. Set the MaxSize to 70% of your heap or queue size. For this example, the value is set to 7. You can get an approximate size of each heap or queue by dividing the number of BaseMap entries by the number of heaps or queues that are used.

SleepTime property

An evictor does not constantly remove entries from your map. Instead it is idle for a set amount of time, only checking the map every n number of seconds, where n refers to the SleepTime property. This property also positively affects performance: running an eviction sweep too often lowers performance because of the resources that are needed for processing them. However, not using the evictor often can result in a map that has entries that are not needed. A map full of entries that are not needed can negatively affect both the memory requirements and processing resources that are required for your map. Setting the eviction sweep interval to fifteen seconds is a good practice for most maps. If the map is written to frequently and is used at a high transaction rate, consider setting the value to a lower time. If the map is accessed infrequently, you can set the time to a higher value.

Example

The following example defines a map, creates a new LFU evictor, sets the evictor properties, and sets the map to use the evictor:

```
//Use ObjectGridManager to create/get the ObjectGrid. Refer to
// the ObjectGridManger section
ObjectGrid objGrid = ObjectGridManager.create.....
BackingMap bMap = objGrid.defineMap("SomeMap");

//Set properties assuming 50,000 map entries
LFUEvictor someEvictor = new LFUEvictor();
```

```
someEvictor.setNumberOfHeaps(5000);
someEvictor.setMaxSize(7);
someEvictor.setSleepTime(15);
bMap.setEvictor(someEvictor);
```

Using the LRU evictor is very similar to using an LFU evictor. An example follows:

```
ObjectGrid objGrid = new ObjectGrid();
BackingMap bMap = objGrid.defineMap("SomeMap");

//Set properties assuming 50,000 map entries
LRUEvictor someEvictor = new LRUEvictor();
someEvictor.setNumberOfLRUQueues(5000);
someEvictor.setMaxSize(7);
someEvictor.setSleepTime(15);
bMap.setEvictor(someEvictor);
```

Notice that only two lines are different from the LFUEvictor example.

Plug-ins for transforming cached objects

Consider transforming cached objects to increase the performance of your cache. You can use the `ObjectTransformer` plug-in when your processor usage is high. Up to 60-70 percent of the total processor time is spent serializing and copying entries. By implementing the `ObjectTransformer` plug-in, you can serialize and deserialize objects with your own implementation. You can use a `CollisionArbiter` plug-in to define how change collisions are handled in your domains.

Developing custom arbiters for multi-master replication

Change collisions might occur if the same records can be changed simultaneously in two places. In a multi-master replication topology, domains detect collisions automatically. When a domain detects a collision, it invokes an arbiter. Typically, collisions are resolved using the default collision arbiter. However, an application can provide a custom collision arbiter.

About this task

If a domain receives a replicated entry that collides with a collision record, the default arbiter uses the changes from the lexicographically lowest named domain. For example, if domain A and B generate a conflict for a record, then the change from domain B is ignored. Domain A keeps its version and the record in domain B is changed to match the record from domain A. Domain names are converted to upper case for comparison.

An alternative option is for the multi-master replication topology to call on a custom collision plug-in to decide the outcome. These instructions outline how to develop a custom collision arbiter and configure a multi-master replication topology to use it.

Procedure

1. Develop a custom collision arbiter and integrate it into your application.

The class should implement the interface:

```
com.ibm.websphere.objectgrid.revision.CollisionArbiter
```

A collision plug-in has three choices for deciding the outcome of a collision. It can choose the local copy or the remote copy or it can provide a revised version of the entry. A domain provides the following information to a custom collision arbiter:

- The existing version of the record
- The collision version of the record
- A Session object that must be used to create the revised version of the collided entry

The plug-in method returns an object indicating its decision. The method invoked by the domain to call the plug-in must return true or false, where false means to ignore the collision – the local version remains unchanged and eXtreme Scale will forget that it ever saw the existing version. The method returns a true value if the method used the provided session to create a new, merged version of the record, reconciling the change.

2. In the `objectgrid.xml` file, specify to use the custom arbiter plug-in.

The id must be "CollisionArbiter."

```
<dg:objectGrid name="revisionGrid" txTimeout="10">
  <dg:bean className="com.you.your_application.
    CustomArbiter" id="CollisionArbiter">
    <dg:property name="property" type="java.lang.String"
      value="propertyValue"/>
  </dg:bean>
</dg:objectGrid>
```

ObjectTransformer plug-in

With the ObjectTransformer plug-in, you can serialize, deserialize, and copy objects in the cache for increased performance.

If you see performance issues with processor usage, add an ObjectTransformer plug-in to each map. If you do not provide an ObjectTransformer plug-in, up to 60-70 percent of the total processor time is spent serializing and copying entries.

Purpose

With the ObjectTransformer plug-in, your applications can provide custom methods for the following operations:

- Serialize or deserialize the key for an entry
- Serialize or deserialize the value for an entry
- Copy a key or value for an entry

If no ObjectTransformer plug-in is provided, you must be able to serialize the keys and values because the ObjectGrid uses a serialize and deserialize sequence to copy the objects. This method is expensive, so use an ObjectTransformer plug-in when performance is critical. The copying occurs when an application looks up an object in a transaction for the first time. You can avoid this copying by setting the copy mode of the Map to NO_COPY or reduce the copying by setting the copy mode to COPY_ON_READ. Optimize the copy operation when needed by the application by providing a custom copy method on this plug-in. Such a plug-in can reduce the copy overhead from 65–70 percent to 2/3 percent of total processor time.

The default `copyKey` and `copyValue` method implementations first attempt to use the clone method, if the method is provided. If no clone method implementation is provided, the implementation defaults to serialization.

Object serialization is also used directly when the eXtreme Scale is running in distributed mode. The LogSequence uses the ObjectTransformer plug-in to help serialize keys and values before transmitting the changes to peers in the

ObjectGrid. You must take care when providing a custom serialization method instead of using the built-in Java developer kit serialization. Object versioning is a complex issue and you might encounter problems with version compatibility if you do not ensure that your custom methods are designed for versioning.

The following list describes how the eXtreme Scale tries to serialize both keys and values:

- If a custom ObjectTransformer plug-in is written and plugged in, eXtreme Scale calls methods in the ObjectTransformer interface to serialize keys and values and get copies of object keys and values.
- If a custom ObjectTransformer plug-in is not used, eXtreme Scale serializes and deserializes values according to the default. If the default plug-in is used, each object is implemented as externalizable or is implemented as serializable.
 - If the object supports the Externalizable interface, the writeExternal method is called. Objects that are implemented as externalizable lead to better performance.
 - If the object does not support the Externalizable interface and does implement the Serializable interface, the object is saved using the ObjectOutputStream method.

Using the ObjectTransformer interface

An ObjectTransformer object must implement the ObjectTransformer interface and follow the common ObjectGrid plug-in conventions.

Two approaches, programmatic configuration and XML configuration, are used to add an ObjectTransformer object into the BackingMap configuration as follows.

Programmatically plug in an ObjectTransformer object

The following code snippet creates the custom ObjectTransformer object and adds it to a BackingMap:

```
ObjectGridManager objectGridManager = ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid myGrid = objectGridManager.createObjectGrid("myGrid", false);
BackingMap backingMap = myGrid.getMap("myMap");
MyObjectTransformer myObjectTransformer = new MyObjectTransformer();
backingMap.setObjectTransformer(myObjectTransformer);
```

XML configuration approach to plug in an ObjectTransformer

Assume that the class name of the ObjectTransformer implementation is the com.company.org.MyObjectTransformer class. This class implements the ObjectTransformer interface. An ObjectTransformer implementation can be configured using the following XML:

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">
  <objectGrids>
    <objectGrid name="myGrid">
      <backingMap name="myMap" pluginCollectionRef="myMap" />
    </objectGrid>
  </objectGrids>

  <backingMapPluginCollections>
    <backingMapPluginCollection id="myMap">
      <bean id="ObjectTransformer" className="com.company.org.MyObjectTransformer" />
    </backingMapPluginCollection>
  </backingMapPluginCollections>
</objectGridConfig>
```

ObjectTransformer usage scenarios

You can use the ObjectTransformer plug-in in the following situations:

- Non-serializable object
- Serializable object but improve serialization performance
- Key or value copy

In the following example, ObjectGrid is used to store the Stock class:

```
/**
 * Stock object for ObjectGrid demo
 *
 *
 */
public class Stock implements Cloneable {
    String ticket;
    double price;
    String company;
    String description;
    int serialNumber;
    long lastTransactionTime;
    /**
     * @return Returns the description.
     */
    public String getDescription() {
        return description;
    }
    /**
     * @param description The description to set.
     */
    public void setDescription(String description) {
        this.description = description;
    }
    /**
     * @return Returns the lastTransactionTime.
     */
    public long getLastTransactionTime() {
        return lastTransactionTime;
    }
    /**
     * @param lastTransactionTime The lastTransactionTime to set.
     */
    public void setLastTransactionTime(long lastTransactionTime) {
        this.lastTransactionTime = lastTransactionTime;
    }
    /**
     * @return Returns the price.
     */
    public double getPrice() {
        return price;
    }
    /**
     * @param price The price to set.
     */
    public void setPrice(double price) {
        this.price = price;
    }
    /**
     * @return Returns the serialNumber.
     */
    public int getSerialNumber() {
        return serialNumber;
    }
    /**
     * @param serialNumber The serialNumber to set.
     */
    public void setSerialNumber(int serialNumber) {
        this.serialNumber = serialNumber;
    }
    /**
     * @return Returns the ticket.
     */
    public String getTicket() {
        return ticket;
    }
    /**
     * @param ticket The ticket to set.
     */
}
```

```

    */
    public void setTicket(String ticket) {
        this.ticket = ticket;
    }
    /**
     * @return Returns the company.
     */
    public String getCompany() {
        return company;
    }
    /**
     * @param company The company to set.
     */
    public void setCompany(String company) {
        this.company = company;
    }
    //clone
    public Object clone() throws CloneNotSupportedException
    {
        return super.clone();
    }
}

```

You can write a custom object transformer class for the Stock class:

```

/**
 * Custom implementation of ObjectGrid ObjectTransformer for stock object
 */
/**
 * public class MyStockObjectTransformer implements ObjectTransformer {
 * / * (non-Javadoc)
 * @see
 * com.ibm.websphere.objectgrid.plugins.ObjectTransformer#serializeKey
 * (java.lang.Object,
 * java.io.ObjectOutputStream)
 */
 public void serializeKey(Object key, ObjectOutputStream stream) throws IOException {
     String ticket= (String) key;
     stream.writeUTF(ticket);
 }

 /* (non-Javadoc)
 * @see com.ibm.websphere.objectgrid.plugins.
 ObjectTransformer#serializeValue(java.lang.Object,
 java.io.ObjectOutputStream)
 */
 public void serializeValue(Object value, ObjectOutputStream stream) throws IOException {
     Stock stock= (Stock) value;
     stream.writeUTF(stock.getTicket());
     stream.writeUTF(stock.getCompany());
     stream.writeUTF(stock.getDescription());
     stream.writeDouble(stock.getPrice());
     stream.writeLong(stock.getLastTransactionTime());
     stream.writeInt(stock.getSerialNumber());
 }

 /* (non-Javadoc)
 * @see com.ibm.websphere.objectgrid.plugins.
 ObjectTransformer#inflateKey(java.io.ObjectInputStream)
 */
 public Object inflateKey(ObjectInputStream stream) throws IOException, ClassNotFoundException {
     String ticket=stream.readUTF();
     return ticket;
 }

 /* (non-Javadoc)
 * @see com.ibm.websphere.objectgrid.plugins.
 ObjectTransformer#inflateValue(java.io.ObjectInputStream)
 */
 public Object inflateValue(ObjectInputStream stream) throws IOException, ClassNotFoundException {
     Stock stock=new Stock();
     stock.setTicket(stream.readUTF());
     stock.setCompany(stream.readUTF());
     stock.setDescription(stream.readUTF());
     stock.setPrice(stream.readDouble());
     stock.setLastTransactionTime(stream.readLong());
     stock.setSerialNumber(stream.readInt());
     return stock;
 }

 /* (non-Javadoc)
 * @see com.ibm.websphere.objectgrid.plugins.
 ObjectTransformer#copyValue(java.lang.Object)
 */
 public Object copyValue(Object value) {
     Stock stock = (Stock) value;

```

```

    try {
        return stock.clone();
    }
    catch (CloneNotSupportedException e)
    {
        // display exception message
    }
}

/* (non-Javadoc)
 * @see com.ibm.websphere.objectgrid.plugins.
 * ObjectTransformer#copyKey(java.lang.Object)
 */
public Object copyKey(Object key) {
    String ticket=(String) key;
    String ticketCopy= new String (ticket);
    return ticketCopy;
}
}

```

Then, plug in this custom MyStockObjectTransformer class into the BackingMap:

```

ObjectGridManager ogf=ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid og = ogf.getObjectGrid("NYSE");
BackingMap bm = og.defineMap("NYSEStocks");
MyStockObjectTransformer ot = new MyStockObjectTransformer();
bm.setObjectTransformer(ot);

```

Serialization performance

WebSphere eXtreme Scale uses multiple Java processes to hold data. These processes serialize the data: That is, they convert the data (which is in the form of Java object instances) to bytes and back to objects again as needed to move the data between client and server processes. Marshalling the data is the most expensive operation and must be addressed by the application developer when designing the schema, configuring the data grid and interacting with the data-access APIs.

The default Java serialization and copy routines are relatively slow and can consume 60 to 70 percent of the processor in a typical setup. The following sections are choices for improving the performance of the serialization.

Write an ObjectTransformer for each BackingMap

An ObjectTransformer can be associated with a BackingMap. Your application can have a class that implements the ObjectTransformer interface and provides implementations for the following operations:

- Copying values
- Serializing and inflating keys to and from streams
- Serializing and inflating values to and from streams

The application does not need to copy keys because keys are considered immutable.

For more information, see Plug-ins for serializing and copying cached objects and ObjectTransformer interface best practices.

Note: The ObjectTransformer is only invoked when the ObjectGrid knows about the data that is being transformed. For example, when DataGrid API agents are used, the agents themselves as well as the agent instance data or data returned from the agent must be optimized using custom serialization techniques. The ObjectTransformer is not invoked for DataGrid API agents.

Using entities

When using the EntityManager API with entities, the ObjectGrid does not store the entity objects directly into the BackingMaps. The EntityManager API converts the

entity object to Tuple objects. See For more information, see the topic on using a loader with entity maps and tuples in the *Programming Guide*. Entity maps are automatically associated with a highly optimized ObjectTransformer. Whenever the ObjectMap API or EntityManager API is used to interact with entity maps, the entity ObjectTransformer is invoked.

Custom serialization

There are some cases when objects must be modified to use custom serialization, such as implementing the java.io.Externalizable interface or by implementing the writeObject and readObject methods for classes implementing the java.io.Serializable interface. Custom serialization techniques should be employed when the objects are serialized using mechanisms other than the ObjectGrid API or EntityManager API methods.

For example, when objects or entities are stored as instance data in a DataGrid API agent or the agent returns objects or entities, those objects are not transformed using an ObjectTransformer. The agent, will however, automatically use the ObjectTransformer when using EntityMixIninterface. See DataGrid agents and entity based Maps for further details.

Byte arrays

When using the ObjectMap or DataGrid APIs, the key and value objects are serialized whenever the client interacts with the data grid and when the objects are replicated. To avoid the overhead of serialization, use byte arrays instead of Java objects. Byte arrays are much cheaper to store in memory since the JDK has less objects to search for during garbage collection and they are can be inflated only when needed. Byte arrays should only be used if you do not need to access the objects using queries or indexes. Since the data is stored as bytes, the data can only be accessed through its key.

WebSphere eXtreme Scale can automatically store data as byte arrays using the CopyMode.COPY_TO_BYTES map configuration option, or it can be handled manually by the client. This option will store the data efficiently in memory and can also automatically inflate the objects within the byte array for use by query and indexes on demand.

ObjectTransformer interface best practices

The ObjectTransformer interface uses callbacks to the application to provide custom implementations of common and expensive operations such as object serialization and deep copies on objects.

Overview

For details about the ObjectTransformer interface, see “ObjectTransformer plug-in” on page 210. From a performance viewpoint, and from the CopyMode method information that is in the CopyMode method best practices topic, eXtreme Scale clearly copies the values for all cases except when NO_COPY mode is used. The default copying mechanism that is employed in eXtreme Scale is serialization, which is known as an expensive operation. The ObjectTransformer interface is used in this situation. The ObjectTransformer interface uses callbacks to the application to provide a custom implementation of common and expensive operations, such as object serialization and deep copies on objects.

An application can provide an implementation of the `ObjectTransformer` interface to a map, and eXtreme Scale then delegates to the methods on this object and relies on the application to provide an optimized version of each method in the interface. The `ObjectTransformer` interface follows:

```
public interface ObjectTransformer {
    void serializeKey(Object key, ObjectOutputStream stream) throws IOException;
    void serializeValue(Object value, ObjectOutputStream stream) throws IOException;
    Object inflateKey(ObjectInputStream stream) throws IOException, ClassNotFoundException;
    Object inflateValue(ObjectInputStream stream) throws IOException, ClassNotFoundException;
    Object copyValue(Object value);
    Object copyKey(Object key);
}
```

You can associate an `ObjectTransformer` interface with a `BackingMap` by using the following example code:

```
ObjectGrid g = ...;
BackingMap bm = g.defineMap("PERSON");
MyObjectTransformer ot = new MyObjectTransformer();
bm.setObjectTransformer(ot);
```

Tune object serialization and inflation

Object serialization is typically the most important performance consideration with eXtreme Scale, which uses the default serializable mechanism if an `ObjectTransformer` plug-in is not supplied by the application. An application can provide implementations of either the `Serializable` `readObject` and `writeObject`, or it can have the objects implement the `Externalizable` interface, which is approximately ten times faster. If the objects in the map cannot be modified, then an application can associate an `ObjectTransformer` interface with the `ObjectMap`. The `serialize` and `inflate` methods are provided to allow the application to provide custom code to optimize these operations, given their large performance impact on the system. The `serialize` method serializes the object to the provided stream. The `inflate` method provides the input stream and expects the application to create the object, inflate it using data in the stream and return the object. Implementations of the `serialize` and `inflate` methods must mirror each other.

Tune deep copy operations

After an application receives an object from an `ObjectMap`, eXtreme Scale performs a deep copy on the object value to ensure that the copy in the `BaseMap` map maintains data integrity. The application can then modify the object value safely. When the transaction commits, the copy of the object value in the `BaseMap` map is updated to the new modified value and the application stops using the value from that point on. You could have copied the object again at the commit phase to make a private copy. However, in this case the performance cost of this action was traded off against requiring the application programmer not to use the value after the transaction commits. The default `ObjectTransformer` attempts to use either a clone or a `serialize` and `inflate` pair to generate a copy. The `serialize` and `inflate` pair is the worst case performance scenario. If profiling reveals that `serialize` and `inflate` is a problem for your application, write an appropriate `clone` method to create a deep copy. If you cannot alter the class, then create a custom `ObjectTransformer` plug-in and implement more efficient `copyValue` and `copyKey` methods.

Plug-ins for versioning and comparing cache objects

Use the `OptimisticCallback` plug-in to customize versioning and comparison operations of cache objects when you are using the optimistic locking strategy.

You can provide a pluggable optimistic callback object that implements the `com.ibm.websphere.objectgrid.plugins.OptimisticCallback` interface. For entity maps, a high performance `OptimisticCallback` plug-in is automatically configured.

Purpose

Use the `OptimisticCallback` interface to provide optimistic comparison operations for the values of a map. An `OptimisticCallback` plug-in is required when you use the optimistic locking strategy. The product provides a default `OptimisticCallback` implementation. However, typically your application must plug in its own implementation of the `OptimisticCallback` interface.

Default implementation

The eXtreme Scale framework provides a default implementation of the `OptimisticCallback` interface that is used if the application does not plug in an application-provided `OptimisticCallback` object. The default implementation always returns the special value of `NULL_OPTIMISTIC_VERSION` as the version object for the value and never updates the version object. This action makes optimistic comparison a "no operation" function. In most cases, you do not want the "no operation" function to occur when you are using the optimistic locking strategy. Your applications must implement the `OptimisticCallback` interface and plug in their own `OptimisticCallback` implementations so that the default implementation is not used. However, at least one scenario exists where the default provided `OptimisticCallback` implementation is useful. Consider the following situation:

- A loader is plugged in for the backing map.
- The loader knows how to perform the optimistic comparison without assistance from an `OptimisticCallback` plug-in.

How can the loader perform optimistic versioning without assistance from an `OptimisticCallback` object? The loader has knowledge of the value class object and knows which field of the value object is used as an optimistic versioning value. For example, suppose the following interface is used for the value object for the employees map:

```
public interface Employee
{
    // Sequential sequence number used for optimistic versioning.
    public long getSequenceNumber();
    public void setSequenceNumber(long newSequenceNumber);
    // Other get/set methods for other fields of Employee object.
}
```

In this example, the loader knows that it can use the `getSequenceNumber` method to get the current version information for an `Employee` value object. The loader increments the returned value to generate a new version number before it updates the persistent storage with the new `Employee` value. For a Java database connectivity (JDBC) loader, the current sequence number in the `WHERE` clause of an overqualified SQL `UPDATE` statement is used, and it uses the new generated sequence number to set the sequence number column to the new sequence number value. Another possibility is that the loader makes use of some backend-provided function that automatically updates a hidden column that can be used for optimistic versioning.

In some situations, a stored procedure or trigger can possibly be used to help maintain a column that holds versioning information. If the loader is using one of these techniques for maintaining optimistic versioning information, then the application does not need to provide an `OptimisticCallback` implementation. The

default `OptimisticCallback` implementation is usable in this scenario because the loader can handle optimistic versioning without any assistance from an `OptimisticCallback` object.

Default implementation for entities

Entities are stored in the `ObjectGrid` using tuple objects. The default `OptimisticCallback` implementation behavior is similar to the behavior for non-entity maps. However, the version field in the entity is identified using the `@Version` annotation or the version attribute in the entity descriptor XML file.

The version attribute can be of the following types: `int`, `Integer`, `short`, `Short`, `long`, `Long` or `java.sql.Timestamp`. An entity must only have one version attribute defined. Only set the version attribute during construction. After the entity is persisted, the value of the version attribute should not be modified.

If a version attribute is not configured and the optimistic locking strategy is used, then the entire tuple is implicitly versioned using the entire state of the tuple, which is much more expensive

In the following example, the `Employee` entity has a long version attribute named `SequenceNumber`:

```
@Entity
public class Employee
{
    private long sequence;
    // Sequential sequence number used for optimistic versioning.
    @Version
    public long getSequenceNumber() {
        return sequence;
    }
    public void setSequenceNumber(long newSequenceNumber) {
        this.sequence = newSequenceNumber;
    }
    // Other get/set methods for other fields of Employee object.
}
```

Writing an `OptimisticCallback` plug-in

An `OptimisticCallback` plug-in must to implement the `OptimisticCallback` interface and follow the common `ObjectGrid` plug-in conventions. See the `OptimisticCallback` interface in the API documentation for more information.

The following list provides a description or consideration for each of the methods in the `OptimisticCallback` interface:

NULL_OPTIMISTIC_VERSION

This special value is returned by the `getVersionedObjectForValue` method if the `OptimisticCallback` implementation does not require version checking. The built-in plugin implementation of the `com.ibm.websphere.objectgrid.plugins.builtins.NoVersioningOptimisticCallback` class uses this value because versioning is disabled when you are specifying this plug-in implementation.

getVersionedObjectForValue method

The `getVersionedObjectForValue` method might return a copy of the value or an attribute of the value that can be used for versioning purposes. This method is called whenever an object is associated with a transaction. When no Loader is plugged into a backing map, the backing map uses this value at commit time to perform an optimistic version comparison. The optimistic version comparison is used by the backing map to ensure that the version has not changed after this transaction first accessed the map entry that was modified by this transaction. If another transaction had already modified the version for this map entry, the version comparison fails and the backing map displays an `OptimisticCollisionException` exception to force the transaction to roll back. If a Loader is plugged in, the backing map does not use the optimistic versioning information. Instead, the Loader is responsible for performing the optimistic versioning comparison and updating the versioning information when necessary. The Loader typically gets the initial versioning object from the `LogElement` passed to the `batchUpdate` method on the loader, which is called when a flush operation occurs or a transaction is committed.

The following code shows the implementation used by the `EmployeeOptimisticCallbackImpl` object:

```
public Object getVersionedObjectForValue(Object value)
{
    if (value == null)
    {
        return null;
    }
    else
    {
        Employee emp = (Employee) value;
        return new Long( emp.getSequenceNumber() );
    }
}
```

As demonstrated in the previous example, the `sequenceNumber` attribute is returned in a `java.lang.Long` object as expected by the Loader, which implies that the same person that wrote the Loader either wrote the `EmployeeOptimisticCallbackImpl` implementation or worked closely with the person that implemented the `EmployeeOptimisticCallbackImpl` - for example, agreed on the value returned by the `getVersionedObjectForValue` method. The default `OptimisticCallback` plug-in returns the special value `NULL_OPTIMISTIC_VERSION` as the version object.

updateVersionedObjectForValue method

This method is called whenever a transaction has updated a value and a new versioned object is needed. If the `getVersionedObjectForValue` method returns an attribute of the value, this method typically updates the attribute value with a new version object. If `getVersionedObjectForValue` method returns a copy of the value, this method typically does not complete any actions. The default `OptimisticCallback` plug-in does not complete any actions with this method because the default implementation of `getVersionedObjectForValue` always returns the special value `NULL_OPTIMISTIC_VERSION` as the version object. The following example shows the implementation used by the `EmployeeOptimisticCallbackImpl` object that is used in the `OptimisticCallback` section:

```

public void updateVersionedObjectForValue(Object value)
{
    if ( value != null )
    {
        Employee emp = (Employee) value;
        long next = emp.getSequenceNumber() + 1;
        emp.updateSequenceNumber( next );
    }
}

```

As demonstrated in the previous example, the `sequenceNumber` attribute increments by one so that the next time the `getVersionedObjectForValue` method is called, the `java.lang.Long` value that is returned has a long value that is the original sequence number value plus one, for example, is the next version value for this employee instance. This example implies that the same person that wrote the Loader either wrote `EmployeeOptimisticCallbackImpl` or worked closely with the person that implemented the `EmployeeOptimisticCallbackImpl`.

serializeVersionedValue method

This method writes the versioned value to the specified stream. Depending on the implementation, the versioned value can be used to identify optimistic update collisions. In some implementations, the versioned value is a copy of the original value. Other implementations might have a sequence number or some other object to indicate the version of the value. Because the actual implementation is unknown, this method is provided to perform the appropriate serialization. The default implementation calls the `writeObject` method.

inflateVersionedValue method

This method takes the serialized version of the versioned value and returns the actual versioned value object. Depending on the implementation, the versioned value can be used to identify optimistic update collisions. In some implementations, the versioned value is a copy of the original value. Other implementations might have a sequence number or some other object to indicate the version of the value. Because the actual implementation is unknown, this method is provided to perform the appropriate deserialization. The default implementation calls the `readObject` method.

Using application-provided OptimisticCallback object

You have two approaches to add an application-provided `OptimisticCallback` object into the `BackingMap` configuration: XML configuration and programmatic configuration.

Programmatically plug in an OptimisticCallback object

The following example demonstrates how an application can programmatically plug in an `OptimisticCallback` object for the employee backing map in the local `grid1` `ObjectGrid` instance:

```

import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
import com.ibm.websphere.objectgrid.ObjectGridManager;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.BackingMap;
ObjectGridManager ogManager = ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid og = ogManager.createObjectGrid( "grid1" );
BackingMap bm = dg.defineMap("employees");
EmployeeOptimisticCallbackImpl cb = new EmployeeOptimisticCallbackImpl();
bm.setOptimisticCallback( cb );

```

XML configuration approach to plug in an OptimisticCallback object

The application can use an XML file to plug in its OptimisticCallback object as shown in the following example:

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">
<objectGrids>
  <objectGrid name="grid1">
    <backingMap name="employees" pluginCollectionRef="employees" lockStrategy="OPTIMISTIC" />
  </objectGrid>
</objectGrids>

<backingMapPluginCollections>
  <backingMapPluginCollection id="employees">
    <bean id="OptimisticCallback" className="com.xyz.EmployeeOptimisticCallbackImpl" />
  </backingMapPluginCollection>
</backingMapPluginCollections>
</objectGridConfig>
```

Plug-ins for providing event listeners

You can use the ObjectGridEventListener and MapEventListener plug-ins to configure notifications for various events in the eXtreme Scale cache. Listener plug-ins are registered with an ObjectGrid or BackingMap instance like other eXtreme Scale plug-ins and add integration and customization points for applications and cache providers.

ObjectGridEventListener plug-in

An ObjectGridEventListener plug-in provides eXtreme Scale life cycle events for the ObjectGrid instance, shards, and transactions. Use the ObjectGridEventListener plug-in to receive notifications when significant events occur on an ObjectGrid. These events include ObjectGrid initialization, the beginning of a transaction, the ending a transaction, and destroying an ObjectGrid. To listen for these events, create a class that implements the ObjectGridEventListener interface and add it to the eXtreme Scale.

For more information about writing an ObjectGridEventListener plug-in, see “ObjectGridEventListener plug-in” on page 223. You can also refer to the API documentation for more information.

Adding and removing ObjectGridEventListener instances

An ObjectGrid can have multiple ObjectGridEventListener listeners. Add and remove the listeners using the addEventListener, setEventListeners and removeEventListener methods on the ObjectGrid interface. You can also declaratively registerObjectGridEventListener plug-ins with the ObjectGrid descriptor file. For examples, see “ObjectGridEventListener plug-in” on page 223.

MapEventListener plug-in

A MapEventListener plug-in provides callback notifications and significant cache state changes that occur for a BackingMap instance. For details on writing a MapEventListener plug-in, see “MapEventListener plug-in” on page 222. You can also refer to the API documentation for more information.

Adding and removing MapEventListener instances

An eXtreme Scale can have multiple MapEventListener listeners. Add and remove listeners with the addMapEventListener, setMapEventListeners and removeMapEventListener methods on the BackingMap interface. You can also declaratively register MapEventListener listeners with the ObjectGrid descriptor file. For examples, see “MapEventListener plug-in.”

MapEventListener plug-in

A MapEventListener plug-in provides callback notifications and significant cache state changes that occur for a BackingMap object: when a map has finished pre-loading or when an entry is evicted from the map. A particular MapEventListener plug-in is a custom class you write implementing the MapEventListener interface.

MapEventListener plug-in conventions

When you develop a MapEventListener plug-in, you must follow common plug-in conventions. For more information about plug-in conventions, see “Introduction to plug-ins” on page 193. For other types of listener plug-ins, see “Plug-ins for providing event listeners” on page 221.

After you write a MapEventListener implementation, you can plug it in to the BackingMap configuration programmatically or with an XML configuration.

Write a MapEventListener implementation

Your application can include an implementation of the MapEventListener plug-in. The plug-in must implement the MapEventListener interface to receive significant events about a map. Events are sent to the MapEventListener plug-in when an entry is evicted from the map and when the preload of a map completes.

Programmatically plug in a MapEventListener implementation

The class name for the custom MapEventListener is the com.company.org.MyMapEventListener class. This class implements the MapEventListener interface. The following code snippet creates the custom MapEventListener object and adds it to a BackingMap object:

```
ObjectGridManager objectGridManager =
    ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid myGrid = objectGridManager.createObjectGrid("myGrid", false);
BackingMap myMap = myGrid.defineMap("myMap");
MyMapEventListener myListener = new MyMapEventListener();
myMap.addMapEventListener(myListener);
```

Plug in a MapEventListener implementation using XML

A MapEventListener implementation can be configured using XML. The following XML must be in the myGrid.xml file:

```
<?xml version="1.0" encoding="UTF-8" ?>
<objectGridconfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/ws/objectgrid/config../objectGrid.xsd"
xmlns="http://ibm.com/ws/objectgrid/config">
  <objectGrids>
    <objectGrid name="myGrid">
      <backingMap name="myMap" pluginCollectionRef="myPlugins" />
    </objectGrid>
  </objectGrids>
  <backingMapPluginCollections>
    <backingMapPluginCollection id="myPlugins">
```



```

        <bean id="MapEventListener" className=
            "com.company.org.MyMapEventListener" />
    </backingMapPluginCollection>
</backingMapPluginCollections>
</objectGridConfig>

```

Providing this file to the ObjectGridManager instance facilitates the creation of this configuration. The following code snippet shows how to create an ObjectGrid instance using this XML file. The newly created ObjectGrid instance has a MapEventListener set on the myMap BackingMap.

```

ObjectGridManager objectGridManager =
    ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid myGrid =
    objectGridManager.createObjectGrid("myGrid", new URL("file:etc/test/myGrid.xml"),
        true, false);

```

ObjectGridEventListener plug-in

An ObjectGridEventListener plug-in provides WebSphere eXtreme Scale life cycle events for the ObjectGrid, shards and transactions. An ObjectGridEventListener plug-in provides notifications when an ObjectGrid is initialized or destroyed, and when a transaction is started or ended. ObjectGridEventListener plug-ins are custom classes you write implementing the ObjectGridEventListener interface. Optionally, the implementation includes ObjectGridEventGroup sub-interfaces and follow the common eXtreme Scale plug-in conventions.

Overview

An ObjectGridEventListener plug-in is useful when a Loader plug-in is available, and you must initialize Java Database Connectivity (JDBC) connections or connections to a back end when transactions start and end. Typically, an ObjectGridEventListener plug-in and a Loader plug-in are written together.

Writing an ObjectGridEventListener plug-in

An ObjectGridEventListener plug-in must implement the ObjectGridEventListener interface to receive notifications about significant eXtreme Scale events. To receive additional event notifications, you can implement the following interfaces. These sub-interfaces are included in the ObjectGridEventGroup interface:

- ShardEvents interface
- ShardLifecycle interface
- TransactionEvents interface

For more information about these interfaces, see the API documentation.

Shard events

When the catalog service places partition primary or replica shards in a Java virtual machine (JVM), a new ObjectGrid instance is created in that JVM to host that shard. Some applications that need to start threads on the JVM host the primary need notification of these events. The ObjectGridEventGroup.ShardEvents interface declares the shardActivate and shardDeactivate methods. These methods are called only when a shard is activated as a primary and when the shard is deactivated from a primary. These two events allow the application to start additional threads when the shard is a primary and stop the threads when the shard returns to being a replica or is taken out of service.

An application can determine which partition has been activated by looking up a specific BackingMap in the ObjectGrid reference that is provided to the shardActivate method using the ObjectGrid#getMap method. The application can then see the partition number using the BackingMap#getPartitionId() method. The partitions are numbered from 0 to the number of partitions in the deployment descriptor minus one.

Shard life-cycle events

ObjectGridEventListener.initialize and ObjectGridEventListener.destroy method events are delivered using the ObjectGridEventGroup.ShardLifecycle interface.

Transaction events

ObjectGridEventListener.transactionBegin and ObjectGridEventListener.transactionEnd methods are delivered through the ObjectGridEventGroup.TransactionEvents interface.

If an ObjectGridEventListener plug-in implements the ObjectGridEventListener and ShardLifecycle interfaces, then shard life-cycle events are the only events that are delivered to the listener. After you implement any of the new ObjectGridEventGroup inner interfaces, eXtreme Scale only delivers those specific events by the new interfaces. With this implementation, code can be backwards compatible. If you are using the new inner interfaces, it can now receive just the specific events that are needed.

Using the ObjectGridEventListener plug-in

To use a custom ObjectGridEventListener plug-in, first create a class that implements the ObjectGridEventListener interface and any optional ObjectGridEventGroup sub-interfaces. Add the custom listener to an ObjectGrid to receive notification of significant events. You have two approaches to add an ObjectGridEventListener plug-in into the eXtreme Scale configuration: programmatic configuration and XML configuration.

Configure an ObjectGridEventListener plug-in programmatically

Assume that the class name of the eXtreme Scale event listener is the com.company.org.MyObjectGridEventListener class. This class implements the ObjectGridEventListener interface. The following code snippet creates the custom ObjectGridEventListener and adds it to an ObjectGrid.

```
ObjectGridManager objectGridManager = ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid myGrid = objectGridManager.createObjectGrid("myGrid", false);
MyObjectGridEventListener myListener = new MyObjectGridEventListener();
myGrid.addEventListener(myListener);
```

Configure an ObjectGridEventListener plug-in with XML

You can also configure an ObjectGridEventListener plug-in using XML. The following XML creates a configuration that is equivalent to the described programmatically created ObjectGrid event listener. The following text must be in the myGrid.xml file:

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">
  <objectGrids>
    <objectGrid name="myGrid">
```

```

        <bean id="ObjectGridEventListener"
            className="com.company.org.MyObjectGridEventListener" />
        <backingMap name="Book"/>
    </objectGrid>
</objectGrids>
</objectGridConfig>

```

Notice the bean declarations come before the backingMap declarations. Provide this file to the ObjectGridManager plug-in to facilitate the creation of this configuration. The following code snippet demonstrates how to create an ObjectGrid instance using this XML file. The ObjectGrid instance that is created has an ObjectGridEventListener listener set on the myGrid ObjectGrid.

```

ObjectGridManager objectGridManager = ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid myGrid = objectGridManager.createObjectGrid("myGrid",
    new URL("file:etc/test/myGrid.xml"), true, false);

```

Plug-ins for custom indexing of cache objects

With a MapIndexPlugin plug-in, or index, you can write custom indexing strategies that are beyond the built-in indexes that eXtreme Scale provides.

For general information about indexing, see [Indexing](#).

For information about using indexing, see “Data access with indexes (Index API)” on page 25.

MapIndexPlugin implementations must use the MapIndexPlugin interface and follow the common eXtreme Scale plug-in conventions.

The following sections include some of the important methods of the index interface.

setProperty method

Use the setProperties method to initialize the index plug-in programmatically. The Properties object parameter that is passed into the method should contain required configuration information to initialize the index plug-in properly. The setProperties method implementation, along with the getProperties method implementation, are required in a distributed environment because the index plug-in configuration moves between client and server processes. An implementation example of this method follows.

```

setProperty(Properties properties)

// setProperties method sample code
public void setProperties(Properties properties) {
    ivIndexProperties = properties;

    String ivRangeIndexString = properties.getProperty("rangeIndex");
    if (ivRangeIndexString != null && ivRangeIndexString.equals("true")) {
        setRangeIndex(true);
    }
    setName(properties.getProperty("indexName"));
    setAttributeName(properties.getProperty("attributeName"));

    String ivFieldAccessAttributeString = properties.getProperty("fieldAccessAttribute");
    if (ivFieldAccessAttributeString != null && ivFieldAccessAttributeString.equals("true")) {
        setFieldAccessAttribute(true);
    }

    String ivPOJOKeyIndexString = properties.getProperty("POJOKeyIndex");
    if (ivPOJOKeyIndexString != null && ivPOJOKeyIndexString.equals("true")) {
        setPOJOKeyIndex(true);
    }
}

```

getProperties method

The `getProperties` method extracts the index plug-in configuration from a `MapIndexPlugin` instance. You can use the extracted properties to initialize another `MapIndexPlugin` instance to have the same internal states. The `getProperties` method and `setProperties` method implementations are required in a distributed environment. An implementation example of the `getProperties` method follows.

```
getProperties()

// getProperties method sample code
public Properties getProperties() {
    Properties p = new Properties();
    p.put("indexName", indexName);
    p.put("attributeName", attributeName);
    p.put("rangeIndex", ivRangeIndex ? "true" : "false");
    p.put("fieldAccessAttribute", ivFieldAccessAttribute ? "true" : "false");
    p.put("POJOKeyIndex", ivPOJOKeyIndex ? "true" : "false");
    return p;
}
```

setEntityMetadata method

The `setEntityMetadata` method is called by the WebSphere eXtreme Scale run time during initialization to set the `EntityMetadata` of the associated `BackingMap` on the `MapIndexPlugin` instance. The `EntityMetadata` is required for supporting indexing of tuple objects. A tuple is a data set that represents an entity object or its key. If the `BackingMap` is for an entity, then you must implement this method.

The following code sample implements the `setEntityMetadata` method.

```
setEntityMetadata(EntityMetadata entityMetadata)

// setEntityMetadata method sample code
public void setEntityMetadata(EntityMetadata entityMetadata) {
    ivEntityMetadata = entityMetadata;
    if (ivEntityMetadata != null) {
        // this is a tuple map
        TupleMetadata valueMetadata = ivEntityMetadata.getValueMetadata();
        int numAttributes = valueMetadata.getNumAttributes();
        for (int i = 0; i < numAttributes; i++) {
            String tupleAttributeName = valueMetadata.getAttribute(i).getName();
            if (attributeName.equals(tupleAttributeName)) {
                ivTupleValueIndex = i;
                break;
            }
        }

        if (ivTupleValueIndex == -1) {
            // did not find the attribute in value tuple, try to find it on key tuple.
            // if found on key tuple, implies key indexing on one of tuple key attributes.
            TupleMetadata keyMetadata = ivEntityMetadata.getKeyMetadata();
            numAttributes = keyMetadata.getNumAttributes();
            for (int i = 0; i < numAttributes; i++) {
                String tupleAttributeName = keyMetadata.getAttribute(i).getName();
                if (attributeName.equals(tupleAttributeName)) {
                    ivTupleValueIndex = i;
                    ivKeyTupleAttributeIndex = true;
                    break;
                }
            }
        }

        if (ivTupleValueIndex == -1) {
            // if entityMetadata is not null and we could not find the
            // attributeName in entityMetadata, this is an
            // error
            throw new ObjectGridRuntimeException("Invalid attributeName. Entity: " +
                ivEntityMetadata.getName());
        }
    }
}
```

Attribute name methods

The `setAttributeName` method sets the name of the attribute to be indexed. The cached object class must provide the `get` method for the indexed attribute. For example, if the object has an `employeeName` or `EmployeeName` attribute, the index calls the `getEmployeeName` method on the object to extract the attribute value. The attribute name must be the same as the name in the `get` method, and the attribute must implement the `Comparable` interface. If the attribute is boolean type, you can also use the `isAttributeName` method pattern.

The `getAttributeName` method returns the name of the indexed attribute.

getAttribute method

The `getAttribute` method returns the indexed attribute value from the specified object. For example, if an `Employee` object has an attribute called `employeeName` that is indexed, you can use the `getAttribute` method to extract the `employeeName` attribute value from a specified `Employee` object. This method is required in a distributed WebSphere eXtreme Scale environment.

```
getAttribute(Object value)

// getAttribute method sample code
public Object getAttribute(Object value) throws ObjectGridRuntimeException {
    if (ivPOJOKeyIndex) {
        // In the POJO key indexing case, no need to get attribute from value object.
        // The key itself is the attribute value used to build the index.
        return null;
    }

    try {
        Object attribute = null;
        if (value != null) {
            // handle Tuple value if ivTupleValueIndex != -1
            if (ivTupleValueIndex == -1) {
                // regular value
                if (ivFieldAccessAttribute) {
                    attribute = this.getAttributeField(value).get(value);
                } else {
                    attribute = getAttributeMethod(value).invoke(value, emptyArray);
                }
            } else {
                // Tuple value
                attribute = extractValueFromTuple(value);
            }
        }
        return attribute;
    } catch (InvocationTargetException e) {
        throw new ObjectGridRuntimeException(
            "Caught unexpected Throwable during index update processing,
            index name = " + indexName + ": " + e,
            e);
    } catch (Throwable t) {
        throw new ObjectGridRuntimeException(
            "Caught unexpected Throwable during index update processing,
            index name = " + indexName + ": " + t,
            t);
    }
}
```

Using a composite index

The composite `HashIndex` improves query performance and avoids expensive map scanning. The feature also provides a convenient way for the `HashIndex` API to find cached objects when search criteria involve many attributes.

Improved performance

A composite `HashIndex` provides a fast and convenient way to search for cached objects with multiple attributes in match-searching criteria. The composite index supports full attribute-match searches, but does not support range searches.

Note: Composite indexes do not support the BETWEEN operator in the ObjectGrid query language because BETWEEN would require range support. The greater than (>) and less than (<) conditionals also do not work because they require range indexes.

A composite index can improve performance of queries if the appropriate composite index is available for the WHERE condition. This means that the composite index has exactly the same attributes as involved in the WHERE condition with full attributes matched.

A query might have many attributes involved in a condition as in the following example.

```
SELECT a FROM Address a WHERE a.city='Rochester' AND a.state='MN' AND a.zipcode='55901'
```

Composite index can improve query performance by avoiding scanning map or joining multiple single-attribute index results. In the example, if a composite index is defined with attributes (city,state,zipcode), the query engine can use the composite index to find the entry with city='Rochester', state='MN', and zipcode='55901'. Without composite index and attribute index on city, state, and zipcode attributes, the query engine will have to scan the map or join multiple single-attribute searches, which usually have expensive overhead. Also, querying for the composite index only supports a full-matched pattern.

Configuring a composite index

You can configure composite indexing in three ways: using XML, programmatically, and with entity annotations only for entity maps.

Programmatic configuration

The programmatic example code below will create the same composite index as the preceding XML.

```
HashIndex mapIndex = new HashIndex();
mapIndex.setName("Address.CityStateZip");
mapIndex.setAttributeName("city,state,zipcode");
mapIndex.setRangeIndex(true);

BackingMap bm = objectGrid.defineMap("mymap");
bm.addMapIndexPlugin(mapIndex);
```

Note that configuring a composite index is the same as configuring a regular index with XML except for the attributeName property value. In a composite index case, the value of attributeName is a comma-delimited list of attributes. For example, the value class Address has 3 attributes: city, state, and zipcode. A composite index can be defined with the attributeName property value as "city,state,zipcode" indicating that city, state, and zipcode are included in the composite index.

Also, note that the composite HashIndexes do not support range lookups and therefore cannot have the RangeIndex property set to true.

Using XML

In order to configure a composite index with XML, include code such as below in the configuration file's backingMapPluginCollections element.

Composite index - XML configuration approach

```
<bean id="MapIndexPlugin" className="com.ibm.websphere.objectgrid.plugins.index.HashIndex">
<property name="Name" type="java.lang.String" value="Address.CityStateZip"/>
<property name="AttributeName" type="java.lang.String" value="city,state,zipcode"/>
</bean>
```

With entity annotations

In the entity map case, annotation approach can be used to define a composite index. You can define a list of `CompositeIndex` within `CompositeIndexes` annotation on the entity class level. The `CompositeIndex` has a name and `attributeNames` property. Each `CompositeIndex` is associated with a `HashIndex` instance applied to the entity's associated `BackingMap`. The `HashIndex` is configured as a non-range index.

```
@Entity
@CompositeIndexes({
    @CompositeIndex(name="CityStateZip", attributeNames="city,state,zipcode"),
    @CompositeIndex(name="lastnameBirthday", attributeNames="lastname,birthday")
})
public class Address {
    @Id int id;
    String street;
    String city;
    String state;
    String zipcode;
    String lastname;
    Date birthday;
}
```

The name property for each composite index must be unique within the entity and `BackingMap`. If the name is not specified, a generated name will be used. The `attributeNames` property is used to populate the `HashIndex` `attributeName` with the comma-delimited list of attributes. The attribute names coincide with the persistent field names when the entities are configured to use field-access, or the property name as defined for the JavaBeans naming conventions for property-access entities. For example: If the attribute name is "street", the property getter method is named `getStreet`.

Performing composite index lookups

After a composite index is configured, an application can use the `findAll(Object)` method of the `MapIndex` interface to perform lookups, as below.

```
Session sess = objectgrid.getSession();
ObjectMap map = sess.getMap("MAP_NAME");
MapIndex codeIndex = (MapIndex) map.getIndex("INDEX_NAME");
Object[] compositeValue = new Object[]{ MapIndex.EMPTY_VALUE,
    "MN", "55901"};
Iterator iter = mapIndex.findAll(compositeValue);
```

The `MapIndex.EMPTY_VALUE` is assigned to the `compositeValue[0]` which indicates that the city attribute is excluded from evaluation. Only objects with state attribute equal to "MN" and zipcode attribute equal to "55901" will be included in the result.

The following queries benefit from the previous composite index configuration:

```
SELECT a FROM Address a WHERE a.city='Rochester' AND a.state='MN' AND
a.zipcode='55901'
```

```
SELECT a FROM Address a WHERE a.state='MN' AND a.zipcode='55901'
```

The query engine will find the appropriate composite index and use it to improve query performance in full attribute-match cases.

In some scenarios, the application might need to define multiple composite indexes with overlapped attributes in order to satisfy all queries with full attributes matched. A disadvantage of increasing the number of indexes is the possible performance overhead on map operations.

Migration and interoperability

The only constraint for the use of a composite index is that an application cannot configure it in a distributed environment with heterogeneous containers. Old and new containers cannot be mixed, since older containers will not recognize a composite index configuration. The composite index is just like the existing regular attribute index, except that the former allows indexing over multiple attributes. When using only the regular attribute index, a mixed-container environment is still viable.

Plug-ins for communicating with persistent stores

With a Loader plug-in, an ObjectGrid map can behave as a memory cache for data that is typically kept in a persistent store on either the same system or some other system. Typically, a database or file system is used as the persistent store. A remote Java virtual machine (JVM) can also be used as the source of data, allowing hub-based caches to be built using ObjectGrid. A loader has the logic for reading and writing data to and from a persistent store.

Loaders are backing map plug-ins that are invoked when changes are made to the backing map or when the backing map is unable to satisfy a data request (a cache miss).

See the information about caching concepts in the *Product Overview* for more information.




Figure 3. Loader

WebSphere eXtreme Scale includes two built-in loaders to integrate with relational database back ends. The Java Persistence API (JPA) loaders use the Object-Relational Mapping (ORM) capabilities of both the OpenJPA and Hibernate implementations of the JPA specification.

Using a loader

To add a loader into the BackingMap configuration, you can use programmatic configuration or XML configuration. A loader has the following relationship with a backing map:

- A backing map can have only one loader.
- A client backing map (near cache) cannot have a loader.
- A loader definition can be applied to multiple backing maps, but each backing map has its own loader instance.

Programmatically plug in a Loader

The following snippet of code demonstrates how to plug an application-provided Loader into the backing map for map1 using the ObjectGrid API:

```
import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
import com.ibm.websphere.objectgrid.ObjectGridManager;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.BackingMap;
ObjectGridManager ogManager = ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid og = ogManager.createObjectGrid( "grid" );
BackingMap bm = og.defineMap( "map1" );
MyLoader loader = new MyLoader();
loader.setDatabaseName("testdb");
loader.setIsolationLevel("read committed");
bm.setLoader( loader );
```

This snippet assumes that the MyLoader class is the application-provided class that implements the com.ibm.websphere.objectgrid.plugins.Loader interface. Because the association of a Loader with a backing map cannot be changed after ObjectGrid is initialized, the code must be run before invoking the initialize method of the ObjectGrid interface that is being called. An IllegalStateException exception occurs on a setLoader method call if it is called after initialization has occurred.

The application-provided Loader can have set properties. In the example, the MyLoader loader is used to read and write data from a table in a relational database. The loader must specify the name of the database and the SQL isolation level. The MyLoader loader has the setDataBaseName and setIsolationLevel methods that allow the application to set these two Loader properties.

XML configuration approach to plug in a Loader

An application-provided Loader can also be plugged in by using an XML file. The following example demonstrates how to plug the MyLoader loader into the map1 backing map with the same database name and isolation level Loader properties. You must specify the className for your loader, the database name and connection details, and the isolation level properties. You can use the same XML structure if you are only using a preloader by specifying the preloader classname instead of a complete loader classname.:

```
<?xml version="1.0" encoding="UTF-8" ?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">
  <objectGrids>
    <objectGrid name="grid">
      <backingMap name="map1" pluginCollectionRef="map1" lockStrategy="OPTIMISTIC" />
    </objectGrid>
  </objectGrids>
  <backingMapPluginCollections>
    <backingMapPluginCollection id="map1">
      <bean id="Loader" className="com.myapplication.MyLoader">
        <property name="dataBaseName"
          type="java.lang.String"
          value="testdb"
          description="database name" />
        <property name="isolationLevel"
          type="java.lang.String"
          value="read committed"
          description="iso level" />
      </bean>
    </backingMapPluginCollection>
  </backingMapPluginCollections>
</objectGridConfig>
```

Writing a loader

You can write your own loader plug-in implementation in your applications, which must follow the common WebSphere eXtreme Scale plug-in conventions.

Including a loader plug-in

The Loader interface has the following definition:

```
public interface Loader
{
    static final SpecialValue KEY_NOT_FOUND;
    List get(TxID txid, List keyList, boolean forUpdate) throws LoaderException;
    void batchUpdate(TxID txid, LogSequence sequence) throws
    LoaderException, OptimisticCollisionException;
    void preloadMap(Session session, BackingMap backingMap) throws LoaderException;
}
```

See the information about Loaders in the *Product Overview* for more information.

get method

The backing map calls the Loader get method to get the values that are associated with a key list that is passed as the keyList argument. The get method is required to return a java.lang.util.List list of values, one value for each key that is in the key list. The first value that is returned in the value list corresponds to the first key in the key list, the second value returned in the value list corresponds to the second key in the key list, and so on. If the loader does not find the value for a key in the key list, the Loader is required to return the special KEY_NOT_FOUND value object that is defined in the Loader interface. Because a backing map can be configured to allow null as a valid value, it is very important for the Loader to return the special KEY_NOT_FOUND object when the Loader cannot find the key. This special value allows the backing map to distinguish between a null value and a value that does not exist because the key was not found. If a backing map does not support null values, a Loader that returns a null value instead of the KEY_NOT_FOUND object for a key that does not exist results in an exception.

The forUpdate argument tells the Loader if the application called a get method on the map or a getForUpdate method on the map. See the ObjectMap interface in the API documentation for more information. The Loader is responsible for implementing a concurrency control policy that controls concurrent access to the persistent store. For example, many relational database management systems support the for update syntax on the SQL select statement that is used to read data from a relational table. The Loader can choose to use the for update syntax on the SQL select statement based on whether boolean true is passed as the argument value for the forUpdate parameter of this method. Typically, the Loader uses the for update syntax only when the pessimistic concurrency control policy is used. For an optimistic concurrency control, the Loader never uses for update syntax on the SQL select statement. The Loader is responsible to decide to use the forUpdate argument based on the concurrency control policy that is being used by the Loader.

For an explanation of the txid parameter, see “Plug-ins for managing transaction life cycle events” on page 250.

batchUpdate method

The batchUpdate method is important on the Loader interface. This method is called whenever the eXtreme Scale needs to apply all the current changes to the Loader. The Loader is given a list of changes for the selected Map. The changes are iterated and applied to the backend. The method receives the current TxID value and the changes to apply. The following sample iterates over the set of changes and batches three Java database connectivity (JDBC) statements, one with insert, another with update, and one with delete.

```
import java.util.Collection;
import java.util.Map;
import java.sql.PreparedStatement;
import java.sql.SQLException;
import com.ibm.websphere.objectgrid.TxID;
import com.ibm.websphere.objectgrid.plugins.Loader;
import com.ibm.websphere.objectgrid.plugins.LoaderException;
import com.ibm.websphere.objectgrid.plugins.LogElement;
import com.ibm.websphere.objectgrid.plugins.LogSequence;

public void batchUpdate(TxID tx, LogSequence sequence) throws LoaderException {
    // Get a SQL connection to use.
    Connection conn = getConnection(tx);
    try {
        // Process the list of changes and build a set of prepared
        // statements for executing a batch update, insert, or delete
        // SQL operation.
        Iterator iter = sequence.getPendingChanges();
        while (iter.hasNext()) {
```

```

        LogElement logElement = (LogElement) iter.next();
        Object key = logElement.getKey();
        Object value = logElement.getCurrentValue();
        switch (logElement.getType().getCode()) {
        case LogElement.CODE_INSERT:
            buildBatchSQLInsert(tx, key, value, conn);
            break;
        case LogElement.CODE_UPDATE:
            buildBatchSQLUpdate(tx, key, value, conn);
            break;
        case LogElement.CODE_DELETE:
            buildBatchSQLDelete(tx, key, conn);
            break;
        }
    }
    // Execute the batch statements that were built by above loop.
    Collection statements = getPreparedStatementCollection(tx, conn);
    iter = statements.iterator();
    while (iter.hasNext()) {
        PreparedStatement pstmt = (PreparedStatement) iter.next();
        pstmt.executeBatch();
    }
} catch (SQLException e) {
    LoaderException ex = new LoaderException(e);
    throw ex;
}
}
}

```

The preceding sample illustrates the high level logic of processing the LogSequence argument, but the details of how a SQL insert, update, or delete statement is built are not illustrated. Some of the key points that are illustrated include:

- The getPendingChanges method is called on the LogSequence argument to obtain an iterator over the list of LogElements that the Loader needs to process.
- The LogElement.getType().getCode() method is used to determine if the LogElement is for a SQL insert, update, or delete operation.
- An SQLException exception is caught and is chained to a LoaderException exception that prints to report that an exception occurred during the batch update.
- JDBC batch update support is used to minimize the number of queries to the backend that must be made.

preloadMap method

During the eXtreme Scale initialization, each backing map that is defined is initialized. If a Loader is plugged into a backing map, the backing map invokes the preloadMap method on the Loader interface to allow the loader to prefetch data from its back end and load the data into the map. The following sample assumes the first 100 rows of an Employee table is read from the database and is loaded into the map. The EmployeeRecord class is an application-provided class that holds the employee data that is read from the employee table.

Note: This sample fetches all the data from database and then insert them into the base map of one partition. In a real-world distributed eXtreme Scale deployment scenario, data should be distributed into all the partitions. Refer to “Developing client-based JPA loaders” on page 283 for more information.

```

import java.sql.PreparedStatement;
import java.sql.SQLException;
import com.ibm.websphere.objectgrid.Session;
import com.ibm.websphere.objectgrid.TxID;
import com.ibm.websphere.objectgrid.plugins.Loader;
import com.ibm.websphere.objectgrid.plugins.LoaderException

public void preloadMap(Session session, BackingMap backingMap) throws LoaderException {
    boolean tranActive = false;
    ResultSet results = null;
    Statement stmt = null;
    Connection conn = null;
    try {
        session.beginNoWriteThrough();
        tranActive = true;
        ObjectMap map = session.getMap(backingMap.getName());
    }
}

```

```

    TxID tx = session.getTxID();
    // Get a auto-commit connection to use that is set to
    // a read committed isolation level.
    conn = getAutoCommitConnection(tx);
    // Preload the Employee Map with EmployeeRecord
    // objects. Read all Employees from table, but
    // limit preload to first 100 rows.
    stmt = conn.createStatement();
    results = stmt.executeQuery(SELECT_ALL);
    int rows = 0;
    while (results.next() && rows < 100) {
        int key = results.getInt(EMPNO_INDEX);
        EmployeeRecord emp = new EmployeeRecord(key);
        emp.setLastName(results.getString(LASTNAME_INDEX));
        emp.setFirstName(results.getString(FIRSTNAME_INDEX));
        emp.setDepartmentName(results.getString(DEPTNAME_INDEX));
        emp.updateSequenceNumber(results.getLong(SEQNO_INDEX));
        emp.setManagerNumber(results.getInt(MGRNO_INDEX));
        map.put(new Integer(key), emp);
        ++rows;
    }
    // Commit the transaction.
    session.commit();
    tranActive = false;
} catch (Throwable t) {
    throw new LoaderException("preload failure: " + t, t);
} finally {
    if (tranActive) {
        try {
            session.rollback();
        } catch (Throwable t2) {
            // Tolerate any rollback failures and
            // allow original Throwable to be thrown.
        }
    }
}
// Be sure to clean up other databases resources here
// as well such a closing statements, result sets, etc.
}
}
}

```

This sample illustrates the following key points:

- The `preloadMap` backing map uses the `Session` object that is passed to it as the session argument.
- The `Session.beginNoWriteThrough` method is used to begin the transaction instead of the `begin` method.
- The `Loader` cannot be called for each `put` operation that occurs in this method for loading the map.
- The `Loader` can map columns of the employee table to a field in the `EmployeeRecord` Java object. The `Loader` catches all throwable exceptions that occur and throws a `LoaderException` exception with the caught throwable exception chained to it.
- The `finally` block ensures that any throwable exception that occurs between the time the `beginNoWriteThrough` method is called and the `commit` method is called cause the `finally` block to roll back the active transaction. This action is critical to ensure that any transaction that has been started by the `preloadMap` method is completed before returning to the caller. The `finally` block is a good place to perform other cleanup actions that might be needed, like closing the Java Database Connectivity (JDBC) connection and other JDBC objects.

The `preloadMap` sample is using a SQL select statement that selects all rows of the table. In your application-provided `Loader`, you might need to set one or more `Loader` properties to control how much of the table needs to be preloaded into the map.

Because the `preloadMap` method is only called one time during the `BackingMap` initialization, it is also a good place to run the one time `Loader` initialization code. Even if a `Loader` chooses not to prefetch data from the backend and load the data into the map, it probably needs to perform some other one time initialization to make other methods of the `Loader` more efficient. The following example illustrates caching the `TransactionCallback` object and `OptimisticCallback` object as instance

variables of the Loader so that the other methods of the Loader do not have to make method calls to get access to these objects. This caching of the ObjectGrid plug-in values can be performed because after the BackingMap is initialized, the TransactionCallback and the OptimisticCallback objects cannot be changed or replaced. It is acceptable to cache these object references as instance variables of the Loader.

```
import com.ibm.websphere.objectgrid.Session;
import com.ibm.websphere.objectgrid.BackingMap;
import com.ibm.websphere.objectgrid.plugins.OptimisticCallback;
import com.ibm.websphere.objectgrid.plugins.TransactionCallback;

// Loader instance variables.
MyTransactionCallback ivTcb; // MyTransactionCallback

// extends TransactionCallback
MyOptimisticCallback ivOcb; // MyOptimisticCallback

// implements OptimisticCallback
// ...
public void preloadMap(Session session, BackingMap backingMap) throws LoaderException
[Replication programming]
    // Cache TransactionCallback and OptimisticCallback objects
    // in instance variables of this Loader.
    ivTcb = (MyTransactionCallback) session.getObjectGrid().getTransactionCallback();
    ivOcb = (MyOptimisticCallback) backingMap.getOptimisticCallback();
    // The remainder of preloadMap code (such as shown in prior example).
}
```

For information about preloading and recoverable preloading as it pertains to replication failover, see the information about replication in the *Product Overview*.

Loaders with entity maps

If the loader is plugged into an entity map, the loader must handle tuple objects. Tuple objects are a special entity data format. The loader must conversion data between tuple and other data formats. For example, the get method returns a list of values that correspond to the set of keys that are passed in to the method. The passed-in keys are in the type of Tuple, says key tuples. Assuming that the loader persists the map with a database using JDBC, the get method must convert each key tuple into a list of attribute values that correspond to the primary key columns of the table that is mapped to the entity map, run the SELECT statement with the WHERE clause that uses converted attribute values as criteria to fetch data from database, and then convert the returned data into value tuples. The get method gets data from the database and converts the data into value tuples for passed-in key tuples, and then returns a list of value tuples correspond to the set of tuple keys that are passed in to the caller. The get method can perform one SELECT statement to fetch all data at one time, or run a SELECT statement for each key tuple. For programming details that show how to use the Loader when the data is store using an entity manager, see “Using a loader with entity maps and tuples” on page 240.

JPA loader programming considerations

A Java Persistence API (JPA) Loader is a loader plug-in implementation that uses JPA to interact with the database. Use the following considerations when you develop an application that uses a JPA loader.

eXtreme Scale entity and JPA entity

You can designate any POJO class as an eXtreme Scale entity using eXtreme Scale entity annotations, XML configuration, or both. You can also designate the same POJO class as a JPA entity using JPA entity annotations, XML configuration, or both.

eXtreme Scale entity: An eXtreme Scale entity represents persistent data that is stored in ObjectGrid maps. An entity object is transformed into a key tuple and a value tuple, which are then stored as key-value pairs in the maps. A tuple is an array of primitive attributes.

JPA entity: A JPA entity represents persistent data that is stored in a relational database automatically using container-managed persistence. The data is persisted in some form of a data storage system in the appropriate form, such as database tuples in a database.

When an eXtreme Scale entity is persisted, its relations are stored in other entity maps. For example, when you are persisting a Consumer entity with a one-to-many relation to a ShippingAddress entity, if cascade-persist is enabled, the ShippingAddress entity is stored in the shippingAddress map in tuple format. If you are persisting a JPA entity, the related JPA entities are also persisted to database tables if cascade-persist is enabled. When a POJO class is designated as both an eXtreme Scale entity and a JPA entity, the data can be persisted to both ObjectGrid entity maps and databases. Common uses follow:

- **Preload scenario:** An entity is loaded from a database using a JPA provider and persists it into ObjectGrid entity maps.
- **Loader scenario:** A Loader implementation is plugged in for the ObjectGrid entity maps so an entity stored in ObjectGrid entity maps can be persisted into or loaded from a database using JPA providers.

It is also common that a POJO class is designated as a JPA entity only. In that case, what is stored in the ObjectGrid maps are the POJO instances, versus the entity tuples in the ObjectGrid entity case.

Application design considerations for entity maps

When you are plugging in a JPALoader interface, the object instances are directly stored in the ObjectGrid maps.

However, you are when plugging in a JPAEntityLoader, the entity class is designated as both an eXtreme Scale entity and a JPA entity. In that case, treat this entity as if it has two persistent stores: the ObjectGrid entity maps and the JPA persistence store. The architecture becomes more complicated than the JPALoader case.

For more information about the JPAEntityLoader plug-in and application design considerations, see the information about the JPAEntityLoader plug-in in the *Administration Guide*. This information can also help if you plan to implement your own loader for the entity maps.

Performance considerations

Ensure that you set the proper eager or lazy fetch type for relationships. For example, a bidirectional one-to-many relationship Consumer to ShippingAddress, with OpenJPA to help explain the performance differences. In this example, a JPA query attempts select o from Consumer o where . . . to do a bulk load, and also load all related ShippingAddress objects. A one-to-many relationship defined in the Consumer class follows:

```

@Entity
public class Consumer implements Serializable {

    @OneToMany(mappedBy="consumer",cascade=CascadeType.ALL, fetch =FetchType.EAGER)
    ArrayList <ShippingAddress> addresses;

```

The many-to-one relation consumer defined in the ShippingAddress class follows:

```

@Entity
public class ShippingAddress implements Serializable{

    @ManyToOne(fetch=FetchType.EAGER)
    Consumer consumer;
}

```

If the fetch types of both relationships are configured as eager, OpenJPA uses N+1+1 queries to get all the Consumer objects and ShippingAddress objects, where N is the number of ShippingAddress objects. However if the ShippingAddress is changed to use lazy fetch type as follows, it only takes two queries to get all the data.

```

@Entity
public class ShippingAddress implements Serializable{

    @ManyToOne(fetch=FetchType.LAZY)
    Consumer consumer;
}

```

Although the query returns the same results, having a lower number of queries significantly decreases interaction with the database, which can increase application performance.

JPAEntityLoader plug-in

The JPAEntityLoader plug-in is a built-in Loader implementation that uses Java Persistence API (JPA) to communicate with the database when you are using the EntityManager API. When you are using the ObjectMap API, use the JPALoader loader.

Loader details

Use the JPALoader plug-in when you are storing data using the ObjectMap API. Use the JPAEntityLoader plug-in when you are storing data using the EntityManager API.

Loaders provide two main functions:

1. **get:** In the get method, the JPAEntityLoader plug-in first calls the `javax.persistence.EntityManager.find(Class entityClass, Object key)` method to find the JPA entity. Then the plug-in projects this JPA entity into entity tuples. During the projection, both the tuple attributes and the association keys are stored in the value tuple. After processing each key, the get method returns a list of entity value tuples.
2. **batchUpdate:** The batchUpdate method takes a LogSequence object that contains a list of LogElement objects. Each LogElement object contains a key tuple and a value tuple. To interact with the JPA provider, you must first find the eXtreme Scale entity based on the key tuple. Based on the LogElement type, you run the following JPA calls:
 - **insert:** `javax.persistence.EntityManager.persist(Object o)`
 - **update:** `javax.persistence.EntityManager.merge(Object o)`

- **remove:** `javax.persistence.EntityManager.remove(Object o)`

A LogElement with type **update** makes the JPAEntityLoader call the `javax.persistence.EntityManager.merge(Object o)` method to merge the entity. However, an **update** type LogElement might be the result of either a `com.ibm.websphere.objectgrid.em.EntityManager.merge(object o)` call or an attribute change of the eXtreme Scale EntityManager managed-instance. See the following example:

```
com.ibm.websphere.objectgrid.em.EntityManager em = og.getSession().getEntityManager();
em.getTransaction().begin();
Consumer c1 = (Consumer) em.find(Consumer.class, c.getConsumerId());
c1.setName("New Name");
em.getTransaction().commit();
```

In this example, an update type LogElement is sent to the JPAEntityLoader of the map consumer. The `javax.persistence.EntityManager.merge(Object o)` method is called to the JPA entity manager instead of an attribute update to the JPA-managed entity. Because of this changed behavior, some limitations exist with using this programming model.

Application design rules

Entities have relationships with other entities. Designing an application with relationships involved and with JPAEntityLoader plugged in requires additional considerations. The application should follow the following four rules, described in the following sections.

Limited relationship depth support

The JPAEntityLoader is only supported when using entities without any relationships or entities with single-level relationships. Relationships with more than one level, such as Company > Department > Employee are not supported.

One loader per map

Using the Consumer-ShippingAddress entity relationships as an example, when you load a consumer with eager fetch enabled, you could load all the related ShippingAddress objects. When you persist or merge a Consumer object, you could persist or merge related ShippingAddress objects if cascade-persist or cascade-merge is enabled.

You cannot plug in a loader for the root entity map which stores the Consumer entity tuples. You must configure a loader for each entity map.

Same cascade type for JPA and eXtreme Scale

Reconsider the scenario where the entity Consumer has a one-to-many relationship with ShippingAddress. You can look at the scenario where cascade-persist is enabled for this relationship. When a Consumer object is persisted into eXtreme Scale, the associated N number of ShippingAddress objects are also persisted into eXtreme Scale.

A persist call of the Consumer object with a cascade-persist relationship to ShippingAddress translates to one `javax.persistence.EntityManager.persist(consumer)` method call and N `javax.persistence.EntityManager.persist(shippingAddress)` method calls by the JPAEntityLoader layer. However, these N extra persist calls to ShippingAddress

objects are unnecessary because of the cascade-persist setting from the JPA provider point of view. To solve this problem, eXtreme Scale provides a new method `isCascaded` on the `LogElement` interface. The `isCascaded` method indicates whether the `LogElement` is a result of an eXtreme Scale `EntityManager` cascade operation. In this example, the `JPAEntityLoader` of the `ShippingAddress` map receives *N* `LogElement` objects because of the cascade persist calls. The `JPAEntityLoader` finds out that the `isCascaded` method returns true and then ignores them without making any JPA calls. Therefore, from a JPA point of view, only one `javax.persistence.EntityManager.persist(consumer)` method call is received.

The same behavior is exhibited if you merge an entity or remove an entity with cascade enabled. The cascaded operations are ignored by the `JPAEntityLoader` plug-in.

The design of the cascade support is to replay the eXtreme Scale `EntityManager` operations to the JPA providers. These operations include `persist`, `merge`, and `remove` operations. To enable cascade support, verify that the cascade setting for the JPA and the eXtreme Scale `EntityManager` are the same.

Use entity update with caution

As previously described, the design of the cascade support is to replay eXtreme Scale `EntityManager` operations to the JPA providers. If your application calls the `ogEM.persist(consumer)` method to the eXtreme Scale `EntityManager`, even the associated `ShippingAddress` objects are persisted because of the cascade-persist setting, and the `JPAEntityLoader` only calls the `jpAEM.persist(consumer)` method to the JPA providers.

However, if your application updates a managed entity, this update translates to a JPA `merge` call by the `JPAEntityLoader` plug-in. In this scenario, support for multiple levels of relationships and key associations is not guaranteed. In this case, the best practice is to use the `javax.persistence.EntityManager.merge(o)` method instead of updating a managed entity.

Using a loader with entity maps and tuples

The entity manager converts all entity objects into tuple objects before they are stored in an WebSphere eXtreme Scale map. Every entity has a key tuple and a value tuple. This key-value pair is stored in the associated eXtreme Scale map for the entity. When you are using an eXtreme Scale map with a loader, the loader must interact with the tuple objects.

eXtreme Scale includes loader plug-ins that simplify integration with relational databases. The Java Persistence API (JPA) Loaders use a Java Persistence API to interact with the database and create the entity objects. The JPA loaders are compatible with eXtreme Scale entities.

Tuples

A tuple contains information about the attributes and associations of an entity. Primitive values are stored using their primitive wrappers. Other supported object types are stored in their native format. Associations to other entities are stored as a collection of key tuple objects that represent the keys of the target entities.

Each attribute or association is stored using a zero-based index. You can retrieve the index of each attribute using the `getAttributePosition` or `getAssociationPosition`

methods. After the position is retrieved, it remains unchanged for the duration of the eXtreme Scale life cycle. The position can change when the eXtreme Scale is restarted. The `setAttribute`, `setAssociation` and `setAssociations` methods are used to update the elements in the tuple.

Attention: When you are creating or updating tuple objects, update every primitive field with a non-null value. Primitive values such as `int` should not be null. If you do not change the value to a default, poor performance issues can result, also affecting fields marked with the `@Version` annotation or `version` attribute in the entity descriptor XML file.

The following example further explains how to process tuples. For more information about defining entities for this example, see the information about the order entity schema, which is in the entity manager tutorial in the *Product Overview*. WebSphere eXtreme Scale is configured for using loaders with each of the entities. Additionally, only the Order entity is taken, and this specific entity has a many-to-one relationship with the Customer entity. The attribute name is `customer`, and it has a one-to-many relationship with the OrderLine entity.

Use the Projector to build Tuple objects automatically from entities. Using the Projector can simplify loaders when you are using an object-relational mapping utility such as Hibernate or JPA.

order.java

```
@Entity
public class Order
{
    @Id String orderNumber;
    java.util.Date date;
    @OneToOne(cascade=CascadeType.PERSIST) Customer customer;
    @OneToMany(cascade=CascadeType.ALL, mappedBy="order") @OrderBy("lineNumber") List<OrderLine> lines;
}
```

customer.java

```
@Entity
public class Customer {
    @Id String id;
    String firstName;
    String surname;
    String address;
    String phoneNumber;
}
```

orderLine.java

```
@Entity
public class OrderLine
{
    @Id @ManyToOne(cascade=CascadeType.PERSIST) Order order;
    @Id int lineNumber;
    @OneToOne(cascade=CascadeType.PERSIST) Item item;
    int quantity;
    double price;
}
```

An `OrderLoader` class that implements the `Loader` interface is shown in the following code. The following example assumes that an associated `TransactionCallback` plug-in is defined.

orderLoader.java

```
public class OrderLoader implements com.ibm.websphere.objectgrid.plugins.Loader {
    private EntityMetadata entityMetadata;
```

```

public void batchUpdate(TxID txid, LogSequence sequence)
    throws LoaderException, OptimisticCollisionException {
    ...
}
public List get(TxID txid, List keyList, boolean forUpdate)
    throws LoaderException {
    ...
}
public void preloadMap(Session session, BackingMap backingMap)
    throws LoaderException {
    this.entityMetaData=backingMap.getEntityMetadadata();
}
}
}

```

The instance variable `entityMetaData` is initialized during the `preLoadMap` method call from the eXtreme Scale. The `entityMetaData` variable is not null if the Map is configured to use entities. Otherwise, the value is null.

batchUpdate method

The `batchUpdate` method provides the ability to know what action the application intended to perform. Based on an insert, update or a delete operation, a connection can be opened to the database and the work performed. Because the key and values are of type `Tuple`, they must be transformed so the values make sense on the SQL statement.

The ORDER table was created with the following Data Definition Language (DDL) definition, as shown in the following code:

```

CREATE TABLE ORDER (ORDERNUMBER VARCHAR(250) NOT NULL, DATE TIMESTAMP, CUSTOMER_ID VARCHAR(250))
ALTER TABLE ORDER ADD CONSTRAINT PK_ORDER PRIMARY KEY (ORDERNUMBER)

```

The following code demonstrates how to convert a `Tuple` to an `Object`:

```

public void batchUpdate(TxID txid, LogSequence sequence)
    throws LoaderException, OptimisticCollisionException {
    Iterator iter = sequence.getPendingChanges();
    while (iter.hasNext()) {
        LogElement logElement = (LogElement) iter.next();
        Object key = logElement.getKey();
        Object value = logElement.getCurrentValue();

        switch (logElement.getType().getCode()) {
            case LogElement.CODE_INSERT:
                1)         if (entityMetaData!=null) {
                    // The order has just one key orderNumber
                    2)         String ORDERNUMBER=(String) getKeyAttribute("orderNumber", (Tuple) key);
                    // Get the value of date
                    3)         java.util.Date unFormattedDate = (java.util.Date) getValueAttribute("date", (Tuple) value);
                    // The values are 2 associations. Lets process customer because
                    // the our table contains customer.id as primary key
                    4)         Object[] keys= getForeignKeyForValueAssociation("customer","id", (Tuple) value);
                            //Order to Customer is M to 1. There can only be 1 key
                    5)         String CUSTOMER_ID=(String)keys[0];
                    // parse variable unFormattedDate and format it for the database as formattedDate
                    6)         String formattedDate = "2007-05-08-14.01.59.780272"; // formatted for DB2
                    // Finally, the following SQL statement to insert the record
                    7) //INSERT INTO ORDER (ORDERNUMBER, DATE, CUSTOMER_ID) VALUES(ORDERNUMBER,formattedDate, CUSTOMER_ID)
                            }
                            break;
                    case LogElement.CODE_UPDATE:
                        break;
                    case LogElement.CODE_DELETE:
                        break;
                }
            }
        }
    }
}
// returns the value to attribute as stored in the key Tuple
private Object getKeyAttribute(String attr, Tuple key) {
    //get key metadata
    TupleMetadata keyMD = entityMetaData.getKeyMetadata();
    //get position of the attribute
    int keyAt = keyMD.getAttributePosition(attr);
    if (keyAt > -1) {
        return key.getAttribute(keyAt);
    } else { // attribute undefined
        throw new IllegalArgumentException("Invalid position index for "+attr);
    }
}

```

```

    }
}
// returns the value to attribute as stored in the value Tuple
private Object getValueAttribute(String attr, Tuple value) {
    //similar to above, except we work with value metadata instead
    TupleMetadata valueMD = entityMetaDatum.getValueMetadata();

    int keyAt = valueMD.getAttributePosition(attr);
    if (keyAt > -1) {
        return value.getAttribute(keyAt);
    } else {
        throw new IllegalArgumentException("Invalid position index for "+attr);
    }
}
// returns an array of keys that refer to association.
private Object[] getForeignKeyForValueAssociation(String attr, String fk_attr, Tuple value) {
    TupleMetadata valueMD = entityMetaDatum.getValueMetadata();
    Object[] ro;

    int customerAssociation = valueMD.getAssociationPosition(attr);
    TupleAssociation tupleAssociation = valueMD.getAssociation(customerAssociation);

    EntityMetadata targetEntityMetadata = tupleAssociation.getTargetEntityMetadata();

    Tuple[] customerKeyTuple = ((Tuple) value).getAssociations(customerAssociation);

    int numberOfKeys = customerKeyTuple.length;
    ro = new Object[numberOfKeys];

    TupleMetadata keyMD = targetEntityMetadata.getKeyMetadata();
    int keyAt = keyMD.getAttributePosition(fk_attr);
    if (keyAt < 0) {
        throw new IllegalArgumentException("Invalid position index for " + attr);
    }
    for (int i = 0; i < numberOfKeys; ++i) {
        ro[i] = customerKeyTuple[i].getAttribute(keyAt);
    }

    return ro;
}
}

```

1. Ensure that entityMetadata is not null, which implies that the key and value cache entries are of type Tuple. From the entityMetadata, Key TupleMetadata is retrieved, which really reflects only the key part of Order metadata.
2. Process the KeyTuple and get the value of Key Attribute orderNumber
3. Process the ValueTuple and get the value of attribute date
4. Process the ValueTuple and get the value of Keys from association customer
5. Extract CUSTOMER_ID. Based on relationship, an Order can only have one customer, we will only have one key. Hence the size of keys is 1. For simplicity, we skipped parsing of date to correct format.
6. Because this is an insert operation, the SQL statement is passed onto the data source connection to complete the insert operation.

Transaction demarcation and access to database is covered in “Writing a loader” on page 232.

get method

If the key is not found in the cache, call the get method in the Loader plug-in to find the key.

The key is a Tuple. The first step is to convert the Tuple to primitive values that can be passed onto the SELECT SQL statement. After all the attributes are retrieved from the database, you must convert into Tuples. The following code demonstrates the Order class.

```

public List get(TxID txid, List keyList, boolean forUpdate) throws LoaderException {
    System.out.println("OrderLoader: Get called");
    ArrayList returnList = new ArrayList();

    1) if (entityMetadata != null) {
        int index=0;
        for (Iterator iter = keyList.iterator(); iter.hasNext();) {

```

```

2) Tuple orderKeyTuple=(Tuple) iter.next();

// The order has just one key orderNumber
3) String ORDERNUMBERKEY = (String) getKeyAttribute("orderNumber",orderKeyTuple);
//We need to run a query to get values of
4) // SELECT CUSTOMER_ID, date FROM ORDER WHERE ORDERNUMBER='ORDERNUMBERKEY'

5) //1) Foreign key: CUSTOMER_ID
6) //2) date
// Assuming those two are returned as
7) String CUSTOMER_ID = "C001"; // Assuming Retrieved and initialized
8) java.util.Date retrievedDate = new java.util.Date();
// Assuming this date reflects the one in database

// We now need to convert this data into a tuple before returning

//create a value tuple
9) TupleMetadata valueMD = entityMetadata.getValueMetadata();
Tuple valueTuple=valueMD.createTuple();

//add retrievedDate object to Tuple
int datePosition = valueMD.getAttributePosition("date");
10) valueTuple.setAttribute(datePosition, retrievedDate);

//Next need to add the Association
11) int customerPosition=valueMD.getAssociationPosition("customer");
TupleAssociation customerTupleAssociation =
valueMD.getAssociation(customerPosition);
EntityMetadata customerEMD = customerTupleAssociation.getTargetEntityMetadata();
TupleMetadata customerTupleMDForKEY=customerEMD.getKeyMetadata();
12) int customerKeyAt=customerTupleMDForKEY.getAttributePosition("id");

Tuple customerKeyTuple=customerTupleMDForKEY.createTuple();
customerKeyTuple.setAttribute(customerKeyAt, CUSTOMER_ID);
13) valueTuple.addAssociationKeys(customerPosition, new Tuple[] {customerKeyTuple});

14) int linesPosition = valueMD.getAssociationPosition("lines");
TupleAssociation linesTupleAssociation = valueMD.getAssociation(linesPosition);
EntityMetadata orderLineEMD = linesTupleAssociation.getTargetEntityMetadata();
TupleMetadata orderLineTupleMDForKEY = orderLineEMD.getKeyMetadata();
int lineNumberAt = orderLineTupleMDForKEY.getAttributePosition("lineNumber");
int orderAt = orderLineTupleMDForKEY.getAssociationPosition("order");

if (lineNumberAt < 0 || orderAt < 0) {
throw new IllegalArgumentException(
"Invalid position index for lineNumber or order "+
lineNumberAt + " " + orderAt);
}
15) // SELECT LINENUMBER FROM ORDERLINE WHERE ORDERNUMBER='ORDERNUMBERKEY'
// Assuming two rows of line number are returned with values 1 and 2

Tuple orderLineKeyTuple1 = orderLineTupleMDForKEY.createTuple();
orderLineKeyTuple1.setAttribute(lineNumberAt, new Integer(1));// set Key
orderLineKeyTuple1.addAssociationKey(orderAt, orderKeyTuple);

Tuple orderLineKeyTuple2 = orderLineTupleMDForKEY.createTuple();
orderLineKeyTuple2.setAttribute(lineNumberAt, new Integer(2));// Init Key
orderLineKeyTuple2.addAssociationKey(orderAt, orderKeyTuple);

16) valueTuple.addAssociationKeys(linesPosition, new Tuple[]
{orderLineKeyTuple1, orderLineKeyTuple2 });

returnList.add(index, valueTuple);

index++;
}
}else {
// does not support tuples
}
return returnList;
}

```

1. The get method is called when the ObjectGrid cache could not find the key and requests the loader to fetch. Test for entityMetadata value and proceed if not null.
2. The keyList contains Tuples.
3. Retrieve value of attribute orderNumber.
4. Run query to retrieve date (value) and customer ID (foreign key).
5. CUSTOMER_ID is a foreign key that must be set in the association tuple.
6. The date is a value and should already be set.

7. Since this example does not perform JDBC calls, CUSTOMER_ID is assumed.
8. Since this example does not perform JDBC calls, date is assumed.
9. Create value Tuple.
10. Set the value of date into the Tuple, based on its position.
11. Order has two associations. Start with the attribute customer which refers to the customer entity. You must have the value of ID to set in the Tuple.
12. Find the position of ID on the customer entity.
13. Set the values of the association keys only.
14. Also, lines is an association that must be set up as a group of association keys, in the same way as you did for customer association.
15. Since you must set up keys for the lineNumber associated with this order, run the SQL to retrieve lineNumber values.
16. Set up the association keys in the valueTuple. This completes the creation of the Tuple that is returned to the BackingMap.

This topic offers the steps create tuples, and a description of the Order entity only. Complete similar steps for the other entities, and the entire process that is tied together with the TransactionCallback plug-in. See “Plug-ins for managing transaction life cycle events” on page 250 for details.

Writing a loader with a replica preload controller

A Loader with a replica preload controller is a Loader that implements the ReplicaPreloadController interface in addition to the Loader interface.

The ReplicaPreloadController interface is designed to provide a way for a replica that becomes the primary shard to know whether the previous primary shard has completed the preload process. If the preload is partially completed, the information to pick up where the previous primary left is provided. With the ReplicaPreloadController interface implementation, a replica that becomes the primary continues the preload process where the previous primary left and continues to finish the overall preload.

In a distributed WebSphere eXtreme Scale environment, a map can have replicas and might preload large volume of data during initialization. The preload is a Loader activity and only occurs in the primary map during initialization. The preload might take a long time to complete if a large volume of data is preloaded. If the primary map has completed large portion of preload data, but is stopped for unknown reason during initialization, a replica becomes the primary. In this situation, the preloaded data that was completed by the previous primary is lost because the new primary normally performs an unconditional preload. With an unconditional preload, the new primary starts the preload process from the beginning and the previous preloaded data is ignored. If you want the new primary to pick up where the previous primary left during preload process, provide a Loader that implements the ReplicaPreloadController interface. For more information see the API documentation.

For information about Loaders, see the information about loaders in the *Product Overview*. If you are interested in writing a regular Loader plug-in, see “Writing a loader” on page 232.

The ReplicaPreloadController interface has the following definition:

```
public interface ReplicaPreloadController
{
    public static final class Status
```

```

    {
        static public final Status PRELOADED_ALREADY =
            new Status(K_PRELOADED_ALREADY);
        static public final Status FULL_PRELOAD_NEEDED =
            new Status(K_FULL_PRELOAD_NEEDED);
        static public final Status PARTIAL_PRELOAD_NEEDED =
            new Status(K_PARTIAL_PRELOAD_NEEDED);
    }

    Status checkPreloadStatus(Session session,
        BackingMap bmap);
}

```

The following sections discuss some of the methods of the Loader and ReplicaPreloadController interface.

checkPreloadStatus method

When a Loader implements ReplicaPreloadController interface, the checkPreloadStatus method is called before the preloadMap method during map initialization. The return status of this method determines if the preloadMap method is called. If this method returns Status#PRELOADED_ALREADY, the preload method is not called. Otherwise, the preload method runs. Because of this behavior, this method should serve as the Loader initialization method. You must initialize the Loader properties in this method. This method must return the correct status, or the preload might not work as expected.

```

public Status checkPreloadStatus(Session session,
    BackingMap backingMap) {
    // When a loader implements ReplicaPreloadController interface,
    // this method will be called before preloadMap method during
    // map initialization. Whether the preloadMap method will be
    // called depends on the returned status of this method. So, this
    // method also serve as Loader's initialization method. This method
    // has to return the right status, otherwise the preload may not
    // work as expected.

    // Note: must initialize this loader instance here.
    ivOptimisticCallback = backingMap.getOptimisticCallback();
    ivBackingMapName = backingMap.getName();
    ivPartitionId = backingMap.getPartitionId();
    ivPartitionManager = backingMap.getPartitionManager();
    ivTransformer = backingMap.getObjectTransformer();
    preloadStatusKey = ivBackingMapName + "_" + ivPartitionId;

    try {
        // get the preloadStatusMap to retrieve preload status that
        // could be set by other JVMs.
        ObjectMap preloadStatusMap = session.getMap(ivPreloadStatusMapName);

        // retrieve last recorded preload data chunk index.
        Integer lastPreloadedDataChunk = (Integer) preloadStatusMap
            .get(preloadStatusKey);

        if (lastPreloadedDataChunk == null) {
            preloadStatus = Status.FULL_PRELOAD_NEEDED;
        } else {
            preloadedLastDataChunkIndex = lastPreloadedDataChunk.intValue();
            if (preloadedLastDataChunkIndex == preloadCompleteMark) {
                preloadStatus = Status.PRELOADED_ALREADY;
            } else {
                preloadStatus = Status.PARTIAL_PRELOAD_NEEDED;
            }
        }

        System.out.println("TupleHeapCacheWithReplicaPreloadControllerLoader.
        checkPreloadStatus()
        -> map = " + ivBackingMapName + ", preloadStatusKey = " + preloadStatusKey
            + ", retrieved lastPreloadedDataChunk = " + lastPreloadedDataChunk + ",
            determined preloadStatus = "
            + getStatusString(preloadStatus));

    } catch (Throwable t) {
        t.printStackTrace();
    }
}

```



```

    }
    return preloadStatus;
}

```

preloadMap method

Running the `preloadMap` method depends on the returned result from `checkPreloadStatus` method. If the `preloadMap` method is called, it typically must retrieve preload status information from designated preload status map and determine how to proceed. Ideally, the `preloadMap` method should know if the preload is partially complete and exactly where to start. During the data preload, the `preloadMap` method should update the preload status on the designated preload status map. The preload status that is stored in the preload status map is retrieved by the `checkPreloadStatus` method when it needs to check the preload status.

```

public void preloadMap(Session session, BackingMap backingMap)
    throws LoaderException {
    EntityMetadata emd = backingMap.getEntityMetadata();
    if (emd != null && tupleHeapPreloadData != null) {
        // The getPreLoadData method is similar to fetching data
        // from database. These data will be push into cache as
        // preload process.
        ivPreloadData = tupleHeapPreloadData.getPreLoadData(emd);

        ivOptimisticCallback = backingMap.getOptimisticCallback();
        ivBackingMapName = backingMap.getName();
        ivPartitionId = backingMap.getPartitionId();
        ivPartitionManager = backingMap.getPartitionManager();
        ivTransformer = backingMap.getObjectTransformer();
        Map preloadMap;

        if (ivPreloadData != null) {
            try {
                ObjectMap map = session.getMap(ivBackingMapName);

                // get the preloadStatusMap to record preload status.
                ObjectMap preloadStatusMap = session.
                    getMap(ivPreloadStatusMapName);

                // Note: when this preloadMap method is invoked, the
                // checkPreloadStatus has been called, Both preloadStatus
                // and preloadedLastDataChunkIndex have been set. And the
                // preloadStatus must be either PARTIAL_PRELOAD_NEEDED
                // or FULL_PRELOAD_NEEDED that will require a preload again.

                // If large amount of data will be preloaded, the data usually
                // is divided into few chunks and the preload process will
                // process each chunk within its own tran. This sample only
                // preload few entries and assuming each entry represent a chunk.
                // so that the preload process an entry in a tran to simulate
                // chunk preloading.

                Set entrySet = ivPreloadData.entrySet();
                preloadMap = new HashMap();
                ivMap = preloadMap;

                // The dataChunkIndex represent the data chunk that is in
                // processing
                int dataChunkIndex = -1;
                boolean shouldRecordPreloadStatus = false;
                int numberOfDataChunk = entrySet.size();
                System.out.println("    numberOfDataChunk to be preloaded = "
                    + numberOfDataChunk);

                Iterator it = entrySet.iterator();
                int whileCounter = 0;
                while (it.hasNext()) {
                    whileCounter++;
                    System.out.println("preloadStatusKey = " + preloadStatusKey
                        + " ,
                    whileCounter = " + whileCounter);

                    dataChunkIndex++;

```

```

        // if the current dataChunkIndex <= preloadedLastDataChunkIndex
        // no need to process, because it has been preloaded by
// other JVM before. only need to process dataChunkIndex
// > preloadedLastDataChunkIndex
        if (dataChunkIndex <= preloadedLastDataChunkIndex) {
            System.out.println("ignore current dataChunkIndex = "
+ dataChunkIndex + " that has been previously
preloaded.");
            continue;
        }

        // Note: This sample simulate data chunk as an entry.
        // each key represent a data chunk for simplicity.
        // If the primary server or shard stopped for unknown
// reason, the preload status that indicates the progress
// of preload should be available in preloadStatusMap. A
// replica that become a primary can get the preload status
// and determine how to preload again.
        // Note: recording preload status should be in the same
// tran as putting data into cache; so that if tran
// rollback or error, the recorded preload status is the
// actual status.

        Map.Entry entry = (Entry) it.next();
        Object key = entry.getKey();
        Object value = entry.getValue();
        boolean tranActive = false;

        System.out.println("processing data chunk. map = " +
this.ivBackingMapName + ", current dataChunkIndex = " +
dataChunkIndex + ", key = " + key);

        try {
            shouldRecordPreloadStatus = false; // re-set to false
            session.beginNoWriteThrough();
            tranActive = true;

            if (ivPartitionManager.getNumOfPartitions() == 1) {
                // if just only 1 partition, no need to deal with
// partition.
                // just push data into cache
                map.put(key, value);
                preloadMap.put(key, value);
                shouldRecordPreloadStatus = true;
            } else if (ivPartitionManager.getPartition(key) ==
ivPartitionId) {
                // if map is partitioned, need to consider the
// partition key only preload data that belongs
// to this partition.
                map.put(key, value);
                preloadMap.put(key, value);
                shouldRecordPreloadStatus = true;
            } else {
                // ignore this entry, because it does not belong to
// this partition.
            }

            if (shouldRecordPreloadStatus) {
                System.out.println("record preload status. map = " +
this.ivBackingMapName + ", preloadStatusKey = " +
preloadStatusKey + ", current dataChunkIndex = "
+ dataChunkIndex);
                if (dataChunkIndex == numberOfDataChunk) {
                    System.out.println("record preload status. map = " +
this.ivBackingMapName + ", preloadStatusKey = " +
preloadStatusKey + ", mark complete = " +
preloadCompleteMark);
                    // means we are at the latest data chunk, if commit
// successfully, record preload complete.
// at this point, the preload is considered to be done
                    // use -99 as special mark for preload complete status.

                    preloadStatusMap.get(preloadStatusKey);

                    // a put follow a get will become update if the get
// return an object, otherwise, it will be insert.
                    preloadStatusMap.put(preloadStatusKey, new
Integer(preloadCompleteMark));
                }
            }
        }
    }
}

```

```

                } else {
                    // record preloaded current dataChunkIndex into
                    // preloadStatusMap a put follow a get will become
                    // update if teh get return an object, otherwise, it
                    // will be insert.
                    preloadStatusMap.get(preloadStatusKey);
                    preloadStatusMap.put(preloadStatusKey, new
Integer(dataChunkIndex));
                }

                session.commit();
                tranActive = false;

                // to simulate preloading large amount of data
                // put this thread into sleep for 30 secs.
                // The real app should NOT put this thread to sleep
                Thread.sleep(10000);

                } catch (Throwable e) {
                    e.printStackTrace();
                    throw new LoaderException("preload failed with
exception: " + e, e);
                } finally {
                    if (tranActive && session != null) {
                        try {
                            session.rollback();
                        } catch (Throwable e1) {
                            // preload ignoring exception from rollback
                        }
                    }
                }
            }
        }

        // at this point, the preload is considered to be done for sure
        // use -99 as special mark for preload complete status.
        // this is a insurance to make sure the complete mark is set.
        // besides, when partitioning, each partition does not know when
        // is its last data chunk. so the following block serves as the
        // overall preload status complete reporting.
        System.out.println("Overall preload status complete -> record
preload status. map = " + this.ivBackingMapName + ",
preloadStatusKey = " + preloadStatusKey + ", mark complete = " +
preloadCompleteMark);
        session.begin();
        preloadStatusMap.get(preloadStatusKey);
        // a put follow a get will become update if teh get return an object,
        // otherwise, it will be insert.
        preloadStatusMap.put(preloadStatusKey, new Integer(preloadCompleteMark));
        session.commit();

        ivMap = preloadMap;
    } catch (Throwable e) {
        e.printStackTrace();
        throw new LoaderException("preload failed with exception: " + e, e);
    }
}
}
}

```

Preload status map

You must use a preload status map to support the `ReplicaPreloadController` interface implementation. The `preloadMap` method should always check the preload status stored in the preload status map first and update the preload status in the preload status map whenever it pushes data into the cache. The `checkPreloadStatus` method can retrieve the preload status from preload status map, determine the preload status, and return the status to its caller. The preload status map should be in the same `mapSet` as other maps that have replica preload controller Loaders.

Plug-ins for managing transaction life cycle events

Use the TransactionCallback plug-in to customize versioning and comparison operations of cache objects when you are using the optimistic locking strategy.

You can provide a pluggable optimistic callback object that implements the `com.ibm.websphere.objectgrid.plugins.OptimisticCallback` interface. For entity maps, a high performance OptimisticCallback plug-in is automatically configured.

Purpose

Use the OptimisticCallback interface to provide optimistic comparison operations for the values of a map. An OptimisticCallback implementation is required when you use the optimistic locking strategy. WebSphere eXtreme Scale provides a default OptimisticCallback implementation. However, usually the application must plug in its own implementation of the OptimisticCallback interface. See the information about locking strategies in the *Product Overview* for more information.

Default implementation

The eXtreme Scale framework provides a default implementation of the OptimisticCallback interface that is used if the application does not plug in an application-provided OptimisticCallback object, as demonstrated in the previous section. The default implementation always returns the special value of `NULL_OPTIMISTIC_VERSION` as the version object for the value and never updates the version object. This action makes optimistic comparison a no operation function. In most cases, you do not want the no operation function to occur when you are using the optimistic locking strategy. Your applications must implement the OptimisticCallback interface and plug in their own OptimisticCallback implementations so that the default implementation is not used. However, at least one scenario exists where the default provided OptimisticCallback implementation is useful. Consider the following situation:

- A loader is plugged for the backing map.
- The loader knows how to perform the optimistic comparison without assistance from an OptimisticCallback plug-in.

How can the loader know how to deal with optimistic versioning without assistance from an OptimisticCallback object? The loader has knowledge of the value class object and knows which field of the value object is used as an optimistic versioning value. For example, suppose the following interface is used for the value object for the employees map:

```
public interface Employee
{
    // Sequential sequence number used for optimistic versioning.
    public long getSequenceNumber();
    public void setSequenceNumber(long newSequenceNumber);
    // Other get/set methods for other fields of Employee object.
}
```

In this case, the loader knows that it can use the `getSequenceNumber` method to get the current version information for an Employee value object. The loader increments the returned value to generate a new version number before updating the persistent storage with the new Employee value. For a Java database connectivity (JDBC) loader, the current sequence number in the where clause of an overqualified SQL update statement is used, and it uses the new generated sequence number to set the sequence number column to the new sequence number value.

Another possibility is that the loader makes use of some backend-provided function that automatically updates a hidden column that can be used for optimistic versioning. In some cases, a stored procedure or trigger can possibly be used to help maintain a column that holds versioning information. If the loader is using one of these techniques for maintaining optimistic versioning information, then the application does not need to provide an `OptimisticCallback` implementation. You can use the default `OptimisticCallback` implementation because the loader is able to handle optimistic versioning without any assistance from an `OptimisticCallback` object.

Default implementation for entities

Entities are stored in the `ObjectGrid` using tuple objects. The default `OptimisticCallback` implementation behaves similarly to the behavior for non-entity maps. However, the version field in the entity is identified using the `@Version` annotation or the version attribute in the entity descriptor XML file.

The version attribute can be of the following types: `int`, `Integer`, `short`, `Short`, `long`, `Long` or `java.sql.Timestamp`. An entity should have only one version attribute defined. The version attribute should be set only during construction. After the entity is persisted, the value of the version attribute should not be modified.

If a version attribute is not configured and the optimistic locking strategy is used, then the entire tuple is implicitly versioned using the entire state of the tuple.

In the following example, the `Employee` entity has a long version attribute named `SequenceNumber`:

```
@Entity
public class Employee
{
    private long sequence;
    // Sequential sequence number used for optimistic versioning.
    @Version
    public long getSequenceNumber() {
        return sequence;
    }
    public void setSequenceNumber(long newSequenceNumber) {
        this.sequence = newSequenceNumber;
    }
    // Other get/set methods for other fields of Employee object.
}
```

Writing an `OptimisticCallback` implementation

An `OptimisticCallback` implementation must implement the `OptimisticCallback` interface and follow the common `ObjectGrid` plug-in conventions

The following list provides a description or consideration for each of the methods in the `OptimisticCallback` interface:

NULL_OPTIMISTIC_VERSION

This special value is returned by `getVersionedObjectForValue` method if the default `OptimisticCallback` implementation is used instead of an application-provided `OptimisticCallback` implementation.

getVersionedObjectForValue method

The `getVersionedObjectForValue` method might return a copy of the value or it might return an attribute of the value that can be used for versioning purposes. This method is called whenever an object is associated with a transaction. When no Loader is set into a backing map, the backing map uses this value at commit time to perform an optimistic version comparison. The optimistic version comparison is used by the backing map to ensure that the version has not changed since this transaction first accessed the map entry that was modified by this transaction. If another transaction had already modified the version for this map entry, the version comparison fails and the backing map displays an `OptimisticCollisionException` exception to force rollback of the transaction. If a Loader is plugged in, the backing map does not use the optimistic versioning information. Instead, the Loader is responsible for performing the optimistic versioning comparison and updating the versioning information when necessary. The Loader typically gets the initial versioning object from the `LogElement` passed to the `batchUpdate` method on the Loader, which is called when a flush operation occurs or a transaction is committed.

The following code shows the implementation used by the `EmployeeOptimisticCallbackImpl` object:

```
public Object getVersionedObjectForValue(Object value)
{
    if (value == null)
    {
        return null;
    }
    else
    {
        Employee emp = (Employee) value;
        return new Long( emp.getSequenceNumber() );
    }
}
```

As demonstrated in the previous example, the `sequenceNumber` attribute is returned in a `java.lang.Long` object as expected by the Loader, which implies that the same person that wrote the Loader either wrote the `EmployeeOptimisticCallbackImpl` implementation or worked closely with the person that implemented the `EmployeeOptimisticCallbackImpl` implementation. For example, these people agreed on the value that is returned by the `getVersionedObjectForValue` method. As previously described, the default `OptimisticCallback` implementation returns the special value `NULL_OPTIMISTIC_VERSION` as the version object.

updateVersionedObjectForValue method

The `updateVersionedObjectForValue` method is called when a transaction has updated a value and a new versioned object is needed. If the `getVersionedObjectForValue` method returns an attribute of the value, this method typically updates the attribute value with a new version object. If the `getVersionedObjectForValue` method returns a copy of the value, this method typically would not update. The default `OptimisticCallback` does not update because the default implementation of the `getVersionedObjectForValue` method always returns the special value `NULL_OPTIMISTIC_VERSION` as the version object. The following example shows the implementation used by the `EmployeeOptimisticCallbackImpl` object that is used in the `OptimisticCallback` section:

```

public void updateVersionedObjectForValue(Object value)
{
    if ( value != null )
    {
        Employee emp = (Employee) value;
        long next = emp.getSequenceNumber() + 1;
        emp.updateSequenceNumber( next );
    }
}

```

As demonstrated in the previous example, the `sequenceNumber` attribute is incremented by one so that the next time the `getVersionedObjectForValue` method is called, the `java.lang.Long` value that is returned has a long value that is the original sequence number value. Plus one, for example, is the next version value for this employee instance. Again, this example implies that the same person that wrote the Loader either wrote `EmployeeOptimisticCallbackImpl` implementation or worked closely with the person that implemented the `EmployeeOptimisticCallbackImpl` implementation.

serializeVersionedValue method

This method writes the versioned value to the specified stream. Depending on the implementation, the versioned value can be used to identify optimistic update collisions. In some implementations, the versioned value is a copy of the original value. Other implementations might have a sequence number or some other object to indicate the version of the value. Because the actual implementation is unknown, this method is provided to perform the proper serialization. The default implementation calls the `writeObject` method.

inflateVersionedValue method

This method takes the serialized version of the versioned value and returns the actual versioned value object. Depending on the implementation, the versioned value can be used to identify optimistic update collisions. In some implementations, the versioned value is a copy of the original value. Other implementations might have a sequence number or some other object to indicate the version of the value. Because the actual implementation is unknown, this method is provided to perform the proper deserialization. The default implementation calls the `readObject` method.

Using an application-provided OptimisticCallback implementation

You have two approaches to add an application-provided `OptimisticCallback` into the `BackingMap` configuration: programmatic configuration and XML configuration.

Programmatically plug in an OptimisticCallback implementation

The following example demonstrates how an application can programmatically plug in an `OptimisticCallback` object for the employee backing map in the `grid1` `ObjectGrid` instance:

```

import com.ibm.websphere.objectgrid.ObjectGridManagerFactory;
import com.ibm.websphere.objectgrid.ObjectGridManager;
import com.ibm.websphere.objectgrid.ObjectGrid;
import com.ibm.websphere.objectgrid.BackingMap;
ObjectGridManager ogManager = ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid og = ogManager.createObjectGrid( "grid1" );
BackingMap bm = dg.defineMap("employees");
EmployeeOptimisticCallbackImpl cb = new EmployeeOptimisticCallbackImpl();
bm.setOptimisticCallback( cb );

```

XML configuration approach to plug in an OptimisticCallback implementation

The EmployeeOptimisticCallbackImpl object in the preceding example must implement the OptimisticCallback interface. The application can also use an XML file to plug in its OptimisticCallback object as shown in the following example:

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">
  <objectGrids>
    <objectGrid name="grid1">
      <backingMap name="employees" pluginCollectionRef="employees" lockStrategy="OPTIMISTIC" />
    </objectGrid>
  </objectGrids>

  <backingMapPluginCollections>
    <backingMapPluginCollection id="employees">
      <bean id="OptimisticCallback" className="com.xyz.EmployeeOptimisticCallbackImpl" />
    </backingMapPluginCollection>
  </backingMapPluginCollections>
</objectGridConfig>
```

Transaction processing overview

WebSphere eXtreme Scale uses transactions as its mechanism for interaction with data.

To interact with data, the thread in your application needs its own session. When the application wants to use the ObjectGrid on a thread, call one of the ObjectGrid.getSession methods to obtain a thread. With the session, the application can work with data that is stored in the ObjectGrid maps.

When an application uses a Session object, the session must be in the context of a transaction. A transaction begins and commits or begins and rolls back using the begin, commit, and rollback methods on the Session object. Applications can also work in auto-commit mode, in which the Session automatically begins and commits a transaction whenever an operation is performed on the map. Auto-commit mode cannot group multiple operations into a single transaction, so it is the slower option if you are creating a batch of multiple operations into a single transaction. However, for transactions that contain only one operation, auto-commit is the faster option.

Introduction to plug-in slots

A plug-in slot is a transactional storage space that is reserved for plug-ins that share transactional context. These slots provide a way for eXtreme Scale plug-ins to communicate with each other, share transactional context, and ensure that transactional resources are used correctly and consistently within a transaction.

A plug-in can store transactional context, such as database connection, Java Message Service (JMS) connection, and so on, in a plug-in slot. The stored transactional context can be retrieved by any plug-in that knows the plug-in slot number, which serves as the key to retrieve transactional context.

Using plug-in slots

Plug-in slots are part of the TxID Interface. See the API documentation for more information about the interface. The slots are entries on an ArrayList array. Plug-ins can reserve an entry in the ArrayList array by calling the ObjectGrid.reserveSlot method and indicating that it wants a slot on all TxID objects. After the slots are reserved, plug-ins can put transactional context into slots of each TxID object and

retrieve the context later. The put and get operations are coordinated by slot numbers that are returned by the `ObjectGrid.reserveSlot` method.

A plug-in typically has a life cycle. Using plug-in slots has to fit into the life cycle of plug-in. Typically, the plug-in must reserve plug-in slots during the initialization stage and obtain slot numbers for each slot. During normal run time, the plug-in puts transactional context into the reserved slot in the `TxID` object at the appropriate point. This appropriate point is typically at the beginning of the transaction. The plug-in or other plug-ins can then get the stored transactional context by the slot number from the `TxID` within the transaction.

The plug-in typically performs a clean up by removing transactional context and the slots. The following snippet of code illustrates how to use plug-in slots in a `TransactionCallback` plug-in:

```
public class DatabaseTransactionCallback implements TransactionCallback {
    int connectionSlot;
    int autoCommitConnectionSlot;
    int psCacheSlot;
    Properties ivProperties = new Properties();

    public void initialize(ObjectGrid objectGrid) throws TransactionCallbackException {
        // In initialization stage, reserve desired plug-in slots by calling the
        // reserveSlot method of ObjectGrid and
        // passing in the designated slot name, TxID.SLOT_NAME.
        // Note: you have to pass in this TxID.SLOT_NAME that is designated
        // for application.
        try {
            // cache the returned slot numbers
            connectionSlot = objectGrid.reserveSlot(TxID.SLOT_NAME);
            psCacheSlot = objectGrid.reserveSlot(TxID.SLOT_NAME);
            autoCommitConnectionSlot = objectGrid.reserveSlot(TxID.SLOT_NAME);
        } catch (Exception e) {
        }
    }

    public void begin(TxID tx) throws TransactionCallbackException {
        // put transactional contexts into the reserved slots at the
        // beginning of the transaction.
        try {
            Connection conn = null;
            conn = DriverManager.getConnection(ivDriverUrl, ivProperties);
            tx.putSlot(connectionSlot, conn);
            conn = DriverManager.getConnection(ivDriverUrl, ivProperties);
            conn.setAutoCommit(true);
            tx.putSlot(autoCommitConnectionSlot, conn);
            tx.putSlot(psCacheSlot, new HashMap());
        } catch (SQLException e) {
            SQLException ex = getLastSQLException(e);
            throw new TransactionCallbackException("unable to get connection", ex);
        }
    }

    public void commit(TxID id) throws TransactionCallbackException {
        // get the stored transactional contexts and use them
        // then, clean up all transactional resources.
        try {
            Connection conn = (Connection) id.getSlot(connectionSlot);
            conn.commit();
            cleanUpSlots(id);
        } catch (SQLException e) {
            SQLException ex = getLastSQLException(e);
            throw new TransactionCallbackException("commit failure", ex);
        }
    }

    void cleanUpSlots(TxID tx) throws TransactionCallbackException {
        closePreparedStatements((Map) tx.getSlot(psCacheSlot));
        closeConnection((Connection) tx.getSlot(connectionSlot));
        closeConnection((Connection) tx.getSlot(autoCommitConnectionSlot));
    }

    /**
     * @param map
     */
    private void closePreparedStatements(Map psCache) {
        try {
            Collection statements = psCache.values();
            Iterator iter = statements.iterator();
            while (iter.hasNext()) {
                PreparedStatement stmt = (PreparedStatement) iter.next();
                stmt.close();
            }
        }
    }
}
```

```

        } catch (Throwable e) {
        }
    }

    /**
     * Close connection and swallow any Throwable that occurs.
     * @param connection
     */
    private void closeConnection(Connection connection) {
        try {
            connection.close();
        } catch (Throwable e1) {
        }
    }

    public void rollback(TxID id) throws TransactionCallbackException
        // get the stored transactional contexts and use them
        // then, clean up all transactional resources.
        try {
            Connection conn = (Connection) id.getSlot(connectionSlot);
            conn.rollback();
            cleanUpSlots(id);
        } catch (SQLException e) {
        }
    }

    public boolean isExternalTransactionActive(Session session) {
        return false;
    }

    // Getter methods for the slot numbers, other plug-in can obtain the slot numbers
    // from these getter methods.

    public int getConnectionSlot() {
        return connectionSlot;
    }
    public int getAutoCommitConnectionSlot() {
        return autoCommitConnectionSlot;
    }
    public int getPreparedStatementSlot() {
        return psCacheSlot;
    }
}

```

The following snippet of code illustrates how a Loader can get the stored transactional context that is put by previous TransactionCallback plug-in example:

```

public class DatabaseLoader implements Loader
{
    DatabaseTransactionCallback tcb;
    public void preloadMap(Session session, BackingMap backingMap) throws LoaderException
    {
        // The preload method is the initialization method of the Loader.
        // Obtain interested plug-in from Session or ObjectGrid instance.
        tcb =
        (DatabaseTransactionCallback)session.getObjectGrid().getTransactionCallback();
    }
    public List get(TxID txid, List keyList, boolean forUpdate) throws LoaderException
    {
        // get the stored transactional contexts that is put by tcb's begin method.
        Connection conn = (Connection)txid.getSlot(tcb.getConnectionSlot());
        // implement get here
        return null;
    }
    public void batchUpdate(TxID txid, LogSequence sequence) throws LoaderException,
        OptimisticCollisionException
    {
        // get the stored transactional contexts that is put by tcb's begin method.
        Connection conn = (Connection)txid.getSlot(tcb.getConnectionSlot());
        // implement batch update here ...
    }
}

```

External transaction managers

Typically, eXtreme Scale transactions begin with the Session.begin method and end with the Session.commit method. However, when an ObjectGrid is embedded, an external transaction coordinator can start and end transactions. In this case, you do not need to call the begin or commit methods.

External transaction coordination

The TransactionCallback plug-in is extended with the `isExternalTransactionActive(Session session)` method that associates the eXtreme Scale session with an external transaction. The method header follows:

```
public synchronized boolean isExternalTransactionActive(Session session)
```

For example, eXtreme Scale can be set up to integrate with WebSphere Application Server and WebSphere Extended Deployment.

Also, eXtreme Scale provides a built in plug-in called the WebSphere “Plug-ins for managing transaction life cycle events” on page 250, which describes how to build the plug-in for WebSphere Application Server environments, but you can adapt the plug-in for other frameworks.

The key to this seamless integration is the exploitation of the ExtendedJTATransaction API in WebSphere Application Server Version 5.x and Version 6.x. However, if you are using WebSphere Application Server Version 6.0.2, you must apply APAR PK07848 to support this method. Use the following sample code to associate an ObjectGrid session with a WebSphere Application Server transaction ID:

```
/**
 * This method is required to associate an objectGrid session with a WebSphere
 * Application Server transaction ID.
 */
Map/**/ localIdToSession;
public synchronized boolean isExternalTransactionActive(Session session)
{
    // remember that this localid means this session is saved for later.
    localIdToSession.put(new Integer(jta.getLocalId()), session);
    return true;
}
```

Retrieve an external transaction

Sometimes you might need to retrieve an external transaction service object for the TransactionCallback plug-in to use. In the WebSphere Application Server server, look up the ExtendedJTATransaction object from its namespace as shown in the following example:

```
public J2EETransactionCallback() {
    super();
    localIdToSession = new HashMap();
    String lookupName="java:comp/websphere/ExtendedJTATransaction";
    try
    {
        InitialContext ic = new InitialContext();
        jta = (ExtendedJTATransaction)ic.lookup(lookupName);
        jta.registerSynchronizationCallback(this);
    }
    catch(NotSupportedException e)
    {
        throw new RuntimeException("Cannot register jta callback", e);
    }
    catch(NamingException e){
        throw new RuntimeException("Cannot get transaction object");
    }
}
```

For other products, you can use a similar approach to retrieve the transaction service object.

Control commit by external callback

The TransactionCallback plug-in must receive an external signal to commit or roll back the eXtreme Scale session. To receive this external signal, use the callback from the external transaction service. Implement the external callback interface and register it with the external transaction service. For example, with WebSphere Application Server, implement the SynchronizationCallback interface, as shown in the following example:

```
public class J2EETransactionCallback implements
com.ibm.websphere.objectgrid.plugins.TransactionCallback, SynchronizationCallback {
    public J2EETransactionCallback() {
        super();
        String lookupName="java:comp/websphere/ExtendedJTATransaction";
        localIdToSession = new HashMap();
        try {
            InitialContext ic = new InitialContext();
            jta = (ExtendedJTATransaction)ic.lookup(lookupName);
            jta.registerSynchronizationCallback(this);
        } catch (NotSupportedException e) {
            throw new RuntimeException("Cannot register jta callback", e);
        }
        catch (NamingException e) {
            throw new RuntimeException("Cannot get transaction object");
        }
    }

    public synchronized void afterCompletion(int localId, byte[] arg1, boolean didCommit) {
        Integer lid = new Integer(localId);
        // find the Session for the localId
        Session session = (Session)localIdToSession.get(lid);
        if (session != null) {
            try {
                // if WebSphere Application Server is committed when
                // hardening the transaction to backingMap.
                // We already did a flush in beforeCompletion
                if (didCommit) {
                    session.commit();
                } else {
                    // otherwise rollback
                    session.rollback();
                }
            } catch (NoActiveTransactionException e) {
                // impossible in theory
            } catch (TransactionException e) {
                // given that we already did a flush, this should not fail
            } finally {
                // always clear the session from the mapping map.
                localIdToSession.remove(lid);
            }
        }
    }

    public synchronized void beforeCompletion(int localId, byte[] arg1) {
        Session session = (Session)localIdToSession.get(new Integer(localId));
        if (session != null) {
            try {
                session.flush();
            } catch (TransactionException e) {
                // WebSphere Application Server does not formally define
                // a way to signal the
                // transaction has failed so do this
                throw new RuntimeException("Cache flush failed", e);
            }
        }
    }
}
```

Use eXtreme Scale APIs with the TransactionCallback plug-in

The TransactionCallback plug-in disables autocommit in eXtreme Scale. The normal usage pattern for an eXtreme Scale follows:

```
Session ogSession = ...;
ObjectMap myMap = ogSession.getMap("MyMap");
ogSession.begin();
MyObject v = myMap.get("key");
v.setAttribute("newValue");
myMap.update("key", v);
ogSession.commit();
```

When this TransactionCallback plug-in is in use, eXtreme Scale assumes that the application uses the eXtreme Scale when a container-managed transaction is present. The previous code snippet changes the following code in this environment:

```
public void myMethod() {
    UserTransaction tx = ...;
    tx.begin();
    Session ogSession = ...;
    ObjectMap myMap = ogSession.getMap("MyMap");
    yObject v = myMap.get("key");
    v.setAttribute("newValue");
    myMap.update("key", v);
    tx.commit();
}
```

The myMethod method is similar to a Web application scenario. The application uses the normal UserTransaction interface to begin, commit, and roll back transactions. The eXtreme Scale automatically begins and commits around the container transaction. If the method is an Enterprise JavaBeans (EJB) method that uses the TX_REQUIRES attribute, then remove the UserTransaction reference and the calls to begin and commit transactions and the method works the same way. In this case, the container is responsible for starting and ending the transaction.

WebSphereTransactionCallback plug-in

When you use the WebSphereTransactionCallback plug-in, enterprise applications that are running in a WebSphere Application Server environment can manage ObjectGrid transactions.

When you are using an ObjectGrid session within a method that is configured to use container-managed transactions, the enterprise container automatically begins, commits or rolls back the ObjectGrid transaction. When you are using Java Transaction API (JTA) UserTransaction objects, the ObjectGrid transaction is managed by the UserTransaction object automatically.

For a detailed discussion of the implementation of this plug-in, see “External transaction managers” on page 256.

Note: The ObjectGrid does not support 2-phase, XA transactions. This plug-in does not enlist the ObjectGrid transaction with the transaction manager. Therefore, if the ObjectGrid fails to commit, any other resources that are managed by the XA transaction do not roll back.

Programmatically plug in the WebSphereTransactionCallback object

You can enable the WebSphereTransactionCallback into the ObjectGrid configuration with programmatic configuration or XML configuration. The following code snippet uses the application to the WebSphereTransactionCallback object and adds it to an ObjectGrid:

```
ObjectGridManager objectGridManager = ObjectGridManagerFactory.getObjectGridManager();
ObjectGrid myGrid = objectGridManager.createObjectGrid("myGrid", false);
WebSphereTransactionCallback wsTxCallback= new WebSphereTransactionCallback ();
myGrid.setTransactionCallback(wsTxCallback);
```

XML configuration approach to plug in the WebSphereTransactionCallback object

The following XML configuration creates the WebSphereTransactionCallback object and adds it to an ObjectGrid. The following text must be in the myGrid.xml file:

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">
  <objectGrids>
    <objectGrid name="myGrid">
      <bean id="TransactionCallback" className=
        "com.ibm.websphere.objectgrid.plugins.builtins.WebSphereTransactionCallback" />
    </objectGrid>
  </objectGrids>
</objectGridConfig>
```

Chapter 5. Programming for administrative tasks

In addition to system application programming interfaces (APIs), WebSphere eXtreme Scale also includes administration APIs that allow applications to monitor and administer servers and clients.

Embedded server API

WebSphere eXtreme Scale includes application programming interfaces (APIs) and system programming interfaces for embedding eXtreme Scale servers and clients within your existing Java applications. The following topic describes the available embedded server APIs.

Instantiating the eXtreme Scale server

You can use several properties to configure the eXtreme Scale server instance, which you can retrieve from the `ServerFactory.getServerProperties` method. The `ServerProperties` object is a singleton, so each call to the `getServerProperties` method retrieves the same instance.

You can create a new server with the following code.

```
Server server = ServerFactory.getInstance();
```

All properties set before the first invocation of `getInstance` are used to initialize the server.

Setting server properties

You can set the server properties until the `ServerFactory.getInstance` is called for the first time. The first call of the `getInstance` method instantiates the eXtreme Scale server, and reads all the configured properties. Setting the properties after creation has no effect. The following example shows how to set properties prior to instantiating a `Server` instance.

```
// Get the server properties associated with this process.
ServerProperties serverProperties = ServerFactory.getServerProperties();

// Set the server name for this process.
serverProperties.setServerName("EmbeddedServerA");

// Set the name of the zone this process is contained in.
serverProperties.setZoneName("EmbeddedZone1");

// Set the end point information required to bootstrap to the catalog service.
serverProperties.setCatalogServiceBootstrap("localhost:2809");

// Set the ORB listener host name to use to bind to.
serverProperties.setListenerHost("host.local.domain");

// Set the ORB listener port to use to bind to.
serverProperties.setListenerPort(9010);

// Turn off all MBeans for this process.
serverProperties.setMBeansEnabled(false);

Server server = ServerFactory.getInstance();
```

Embedding the catalog service

Any JVM setting that is flagged by the `CatalogServerProperties.setCatalogServer` method can host the catalog service for eXtreme Scale. This method indicates to the eXtreme Scale server runtime to instantiate the catalog service when the server is started. The following code shows how to instantiate the eXtreme Scale catalog server:

```
CatalogServerProperties catalogServerProperties =
    ServerFactory.getCatalogProperties();
catalogServerProperties.setCatalogServer(true);

Server server = ServerFactory.getInstance();
```

Embedding the eXtreme Scale container

Issue the `Server.createContainer` method for any JVM to host multiple eXtreme Scale containers. The following code shows how to instantiate an eXtreme Scale container:

```
Server server = ServerFactory.getInstance();
DeploymentPolicy policy = DeploymentPolicyFactory.createDeploymentPolicy(
    new File("META-INF/embeddedDeploymentPolicy.xml").toURI().toURL(),
    new File("META-INF/embeddedObjectGrid.xml").toURI().toURL());
Container container = server.createContainer(policy);
```

Self-contained server process

You can start all the services together, which is useful for development and also practical in production. By starting the services together, a single process does all of the following: Starts the catalog service, starts a set of containers, and runs the client connection logic. Starting the services in this way sorts out programming issues prior to deploying in a distributed environment. The following code shows how to instantiate a self-contained eXtreme Scale server:

```
CatalogServerProperties catalogServerProperties =
    ServerFactory.getCatalogProperties();
catalogServerProperties.setCatalogServer(true);

Server server = ServerFactory.getInstance();
DeploymentPolicy policy = DeploymentPolicyFactory.createDeploymentPolicy(
    new File("META-INF/embeddedDeploymentPolicy.xml").toURI().toURL(),
    new File("META-INF/embeddedObjectGrid.xml").toURI().toURL());
Container container = server.createContainer(policy);
```

Embedding eXtreme Scale in WebSphere Application Server

The configuration for eXtreme Scale is set up automatically when you install eXtreme Scale in a WebSphere Application Server environment. You are not required to set any properties before you access the server to create a container. The following code shows how to instantiate an eXtreme Scale server in WebSphere Application Server:

```
Server server = ServerFactory.getInstance();
DeploymentPolicy policy = DeploymentPolicyFactory.createDeploymentPolicy(
    new File("META-INF/embeddedDeploymentPolicy.xml").toURI().toURL(),
    new File("META-INF/embeddedObjectGrid.xml").toURI().toURL());
Container container = server.createContainer(policy);
```

For a step by step example on how to start an embedded catalog service and container programmatically, see “Using the embedded server API to start and stop servers” on page 263.

Using the embedded server API to start and stop servers

With WebSphere eXtreme Scale, you can use a programmatic API for managing the life cycle of embedded servers and containers. You can programmatically configure the server with any of the options that you can also configure with the command line options or file-based server properties. You can configure the embedded server to be a container server, a catalog service, or both.

Before you begin

You must have a method for running code from within an already existing Java virtual machine. The eXtreme Scale classes must be available through the class loader tree.

About this task

You can run many administration tasks with the Administration API. One common use of the API is as an internal server for storing Web application state. The Web server can start an embedded WebSphere eXtreme Scale server, report the container server to the catalog service, and the server is then added as a member of a larger distributed grid. This usage can provide scalability and high availability to an otherwise volatile data store.

You can programmatically control the complete life cycle of an embedded eXtreme Scale server. The examples are as generic as possible and only show direct code examples for the outlined steps.

Procedure

1. Obtain the `ServerProperties` object from the `ServerFactory` class and configure any necessary options.

Every eXtreme Scale server has a set of configurable properties. When a server starts from the command line, those properties are set to defaults, but you can override several properties by providing an external source or file. In the embedded scope, you can directly set the properties with a `ServerProperties` object. You must set these properties before you obtain a server instance from the `ServerFactory` class. The following example snippet obtains a `ServerProperties` object, sets the `CatalogServiceBootstrap` field, and initializes several optional server settings. See the API documentation for a list of the configurable settings.

```
ServerProperties props = ServerFactory.getServerProperties();
props.setCatalogServiceBootstrap("host:port"); // required to connect to specific catalog service
props.setServerName("ServerOne"); // name server
props.setTraceSpecification("com.ibm.ws.objectgrid=all=enabled"); // Sets trace spec
```

2. If you want the server to be a catalog service, obtain the `CatalogServerProperties` object.

Every embedded server can be a catalog service, a container server, or both a container server and a catalog service. The following example obtains the `CatalogServerProperties` object, enables the catalog service option, and configures various catalog service settings.

```
CatalogServerProperties catalogProps = ServerFactory.getCatalogProperties();
catalogProps.setCatalogServer(true); // false by default, it is required to set as a catalog service
catalogProps.setQuorum(true); // enables / disables quorum
```

3. Obtain a `Server` instance from the `ServerFactory` class. The `Server` instance is a process-scoped singleton that is responsible for managing the membership in the grid. After this instance has been instantiated, this process is connected

and is highly available with the other servers in the grid. The following example shows how to create the Server instance:

```
Server server = ServerFactory.getInstance();
```

Reviewing the previous example, the ServerFactory class provides a static method that returns a Server instance. The ServerFactory class is intended to be the only interface for obtaining a Server instance. Therefore, the class ensures that the instance is a singleton, or one instance for each JVM or isolated classloader. The getInstance method initializes the Server instance. You must configure all the server properties before you initialize the instance. The Server class is responsible for creating new Container instances. You can use both the ServerFactory and Server classes for managing the life cycle of the embedded Server instance.

4. Start a Container instance using the Server instance.

Before shards can be placed on an embedded server, you must create a container on the server. The Server interface has a createContainer method that takes a DeploymentPolicy argument. The following example uses the server instance that you obtained to create a container using a created DeploymentPolicy file. Note that Containers require a classloader that has the application binaries available to it for serialization. You can make these binaries available by calling the createContainer method with the Thread context classloader set to the classloader that you want to use.

```
DeploymentPolicy policy = DeploymentPolicyFactory.createDeploymentPolicy(new
    URL("file://urltodeployment.xml"),
    new URL("file://urltoobjectgrid.xml"));
Container container = server.createContainer(policy);
```

5. Remove and clean up a container.

You can remove and clean up a container server by using the running the teardown method on the obtained Container instance. Running the teardown method on a container properly cleans up the container and removes the container from the embedded server.

The process of cleaning up the container includes the movement and tearing down of all the shards that are placed within that container. Each server can contain many containers and shards. Cleaning up a container does not affect the life cycle of the parent Server instance. The following example demonstrates how to run the teardown method on a server. The teardown method is made available through the ContainerMBean interface. By using the ContainerMBean interface, if you no longer have programmatic access to this container, you can still remove and clean up the container with its MBean. A terminate method also exists on the Container interface, do not use this method unless it is absolutely needed. This method is more forceful and does not coordinate appropriate shard movement and clean up.

```
container.teardown();
```

6. Stop the embedded server.

When you stop an embedded server, you also stop any containers and shards that are running on the server. When you stop an embedded server, you must clean up all open connections and move or tear down all the shards. The following example demonstrates how to stop a server and using the waitFor method on the Server interface to ensure that the Server instance shuts down completely. Similarly to the container example, the stopServer method is made available through the ServerMBean interface. With this interface, you can stop a server with the corresponding Managed Bean (MBean).

```
ServerFactory.stopServer(); // Uses the factory to kill the Server singleton
// or
server.stopServer(); // Uses the Server instance directly
server.waitFor(); // Returns when the server has properly completed its shutdown procedures
```

Full code example:

```
import java.net.MalformedURLException;
import java.net.URL;

import com.ibm.websphere.objectgrid.ObjectGridException;
import com.ibm.websphere.objectgrid.deployment.DeploymentPolicy;
import com.ibm.websphere.objectgrid.deployment.DeploymentPolicyFactory;
import com.ibm.websphere.objectgrid.server.Container;
import com.ibm.websphere.objectgrid.server.Server;
import com.ibm.websphere.objectgrid.server.ServerFactory;
import com.ibm.websphere.objectgrid.server.ServerProperties;

public class ServerFactoryTest {

    public static void main(String[] args) {

        try {

            ServerProperties props = ServerFactory.getServerProperties();
            props.setCatalogServiceBootstrap("catalogservice-hostname:catalogservice-port");
            props.setServerName("ServerOne"); // name server
            props.setTraceSpecification("com.ibm.ws.objectgrid=all=enabled"); // TraceSpec

            /*
             * In most cases, the server will serve as a container server only
             * and will connect to an external catalog service. This is a more
             * highly available way of doing things. The commented code excerpt
             * below will enable this Server to be a catalog service.
             *
             *
             * CatalogServerProperties catalogProps =
             * ServerFactory.getCatalogProperties();
             * catalogProps.setCatalogServer(true); // enable catalog service
             * catalogProps.setQuorum(true); // enable quorum
             */

            Server server = ServerFactory.getInstance();

            DeploymentPolicy policy = DeploymentPolicyFactory.createDeploymentPolicy
(new URL("url to deployment xml"), new URL("url to objectgrid xml file"));
            Container container = server.createContainer(policy);

            /*
             * Shard will now be placed on this container if the deployment requirements are met.
             * This encompasses embedded server and container creation.
             *
             * The lines below will simply demonstrate calling the cleanup methods
             */

            container.teardown();
            server.stopServer();
            int success = server.waitFor();

        } catch (ObjectGridException e) {
            // Container failed to initialize
        } catch (MalformedURLException e2) {
            // invalid url to xml file(s)
        }

    }

}
```

Monitoring with the statistics API

The Statistics API is the direct interface to the internal statistics tree. Statistics are disabled by default, but can be enabled by setting a StatsSpec interface. A StatsSpec interface defines how WebSphere eXtreme Scale should monitor statistics.

About this task

You can use the local StatsAccessor API to query data and access statistics on any ObjectGrid instance that is in the same Java virtual machine (JVM) as the running code. For more information about the specific interfaces, see the API documentation. Use the following steps to enable monitoring of the internal statistics tree.

Procedure

1. Retrieve the StatsAccessor object. The StatsAccessor interface follows the singleton pattern. So, apart from problems related to the classloader, one StatsAccessor instance should exist for each JVM. This class serves as the main interface for all local statistics operations. The following code is an example of how to retrieve the accessor class. Call this operation before any other ObjectGrid calls.

```
public class LocalClient
{
    public static void main(String[] args) {
        // retrieve a handle to the StatsAccessor
        StatsAccessor accessor = StatsAccessorFactory.getStatsAccessor();
    }
}
```

2. Set the data grid StatsSpec interface. Set this JVM to collect all statistics at the ObjectGrid level only. You must ensure that an application enables all statistics that might be needed before you begin any transactions. The following example sets the StatsSpec interface using both a static constant field and using a spec String. Using a static constant field is simpler because the field has already defined the specification. However, by using a spec String, you can enable any combination of statistics that are required.

```
public static void main(String[] args) {
    // retrieve a handle to the StatsAccessor
    StatsAccessor accessor = StatsAccessorFactory.getStatsAccessor();

    // Set the spec via the static field
    StatsSpec spec = new StatsSpec(StatsSpec.OG_ALL);
    accessor.setStatsSpec(spec);

    // Set the spec via the spec String
    StatsSpec spec = new StatsSpec("og.all=enabled");
    accessor.setStatsSpec(spec);
}
```

3. Send transactions to the grid to force data to be collected for monitoring. To collect useful data for statistics, you must send transactions to the data grid. The following code excerpt inserts a record into MapA, which is in ObjectGridA. Because the statistics are at the ObjectGrid level, any map within the ObjectGrid yields the same results.

```
public static void main(String[] args) {
    // retrieve a handle to the StatsAccessor
    StatsAccessor accessor = StatsAccessorFactory.getStatsAccessor();

    // Set the spec via the static field
    StatsSpec spec = new StatsSpec(StatsSpec.OG_ALL);
    accessor.setStatsSpec(spec);

    ObjectGridManager manager =
    ObjectGridmanagerFactory.getObjectGridManager();
    ObjectGrid grid = manager.getObjectGrid("ObjectGridA");
    Session session = grid.getSession();
    Map map = session.getMap("MapA");

    // Drive insert
}
```

```

        session.begin();
        map.insert("SomeKey", "SomeValue");
        session.commit();
    }

```

4. Query a StatsFact by using the StatsAccessor API. Every statistics path is associated with a StatsFact interface. The StatsFact interface is a generic placeholder that is used to organize and contain a StatsModule object. Before you can access the actual statistics module, the StatsFact object must be retrieved.

```

public static void main(String[] args)
{
    // retrieve a handle to the StatsAccessor
    StatsAccessor accessor = StatsAccessorFactory.getStatsAccessor();

    // Set the spec via the static field
    StatsSpec spec = new StatsSpec(StatsSpec.OG_ALL);
    accessor.setStatsSpec(spec);

    ObjectGridManager manager =
    ObjectGridManagerFactory.getObjectGridManager();
    ObjectGrid grid = manager.getObjectGrid("ObjectGridA");
    Session session = grid.getSession();
    Map map = session.getMap("MapA");

    // Drive insert
    session.begin();
    map.insert("SomeKey", "SomeValue");
    session.commit();

    // Retrieve StatsFact

    StatsFact fact = accessor.getStatsFact(new String[] {"EmployeeGrid"},
    StatsModule.MODULE_TYPE_OBJECT_GRID);
}

```

5. Interact with the StatsModule object. The StatsModule object is contained within the StatsFact interface. You can obtain a reference to the module by using the StatsFact interface. Since the StatsFact interface is a generic interface, you must cast the returned module to the expected StatsModule type. Because this task collects eXtreme Scale statistics, the returned StatsModule object is cast to an OGStatsModule type. After the module is cast, you have access to all of the available statistics.

```

public static void main(String[] args) {
    // retrieve a handle to the StatsAccessor
    StatsAccessor accessor = StatsAccessorFactory.getStatsAccessor();

    // Set the spec via the static field
    StatsSpec spec = new StatsSpec(StatsSpec.OG_ALL);
    accessor.setStatsSpec(spec);

    ObjectGridManager manager =
    ObjectGridmanagerFactory.getObjectGridManager();
    ObjectGrid grid = manager.getObjectGrid("ObjectGridA");
    Session session = grid.getSession();
    Map map = session.getMap("MapA");

    // Drive insert
    session.begin();
    map.insert("SomeKey", "SomeValue");
    session.commit();
}

```

```

// Retrieve StatsFact
StatsFact fact = accessor.getStatsFact(new String[] {"EmployeeGrid"},
StatsModule.MODULE_TYPE_OBJECT_GRID);

// Retrieve module and time
OGStatsModule module = (OGStatsModule)fact.getStatsModule();
ActiveTimeStatistic timeStat =
module.getTransactionTime("Default", true);
double time = timeStat.getMeanTime();
}

```

Monitoring with WebSphere Application Server PMI

WebSphere eXtreme Scale supports Performance Monitoring Infrastructure (PMI) when running in a WebSphere Application Server or WebSphere Extended Deployment application server. PMI collects performance data on runtime applications and provides interfaces that support external applications to monitor performance data. You can use the administrative console or the wsadmin tool to access monitoring data.

Before you begin

You can use PMI to monitor your environment when you are using WebSphere eXtreme Scale combined with WebSphere Application Server.

About this task

WebSphere eXtreme Scale uses the custom PMI feature of WebSphere Application Server to add its own PMI instrumentation. With this approach, you can enable and disable WebSphere eXtreme Scale PMI with the administrative console or with Java Management Extensions (JMX) interfaces in the wsadmin tool. In addition, you can access WebSphere eXtreme Scale statistics with the standard PMI and JMX interfaces that are used by monitoring tools, including the Tivoli® Performance Viewer.

Procedure

1. Enable eXtreme Scale PMI. You must enable PMI to view the PMI statistics. See “Enabling PMI” for more information.
2. Retrieve eXtreme Scale PMI statistics. View the performance of your eXtreme Scale applications with the Tivoli Performance Viewer. See “Retrieving PMI statistics” on page 270 for more information.

Enabling PMI

You can use WebSphere Application Server Performance Monitoring Infrastructure (PMI) to enable or disable statistics at any level. For example, you can choose to enable the map hit rate statistic for a particular map, but not the number of entry statistic or the loader batch update time statistic. You can enable PMI in the administrative console or with scripting.

Before you begin

Your application server must be started and have an eXtreme Scale-enabled application installed. To enable PMI with scripting, you also must be able to log in and use the wsadmin tool. For more information about the wsadmin tool, see the wsadmin tool topic in the WebSphere Application Server information center.

About this task

Use WebSphere Application Server PMI to provide a granular mechanism with which you can enable or disable statistics at any level. For example, you can choose to enable the map hit rate statistic for a particular map, but not the number of entry or the loader batch update time statistics. This section shows how to use the administrative console and wsadmin scripts to enable ObjectGrid PMI.

Procedure

- **Enable PMI in the administrative console.**

1. In the administrative console, click **Monitoring and Tuning > Performance Monitoring Infrastructure > server_name**.
2. Verify that Enable Performance Monitoring Infrastructure (PMI) is selected. This setting is enabled by default. If the setting is not enabled, select the check box, then restart the server.
3. Click **Custom**. In the configuration tree, select the ObjectGrid and ObjectGrid Maps module. Enable the statistics for each module.

The transaction type category for ObjectGrid statistics is created at runtime. You can see only the subcategories of the ObjectGrid and Map statistics on the **Runtime** tab.

- **Enable PMI with scripting.**

1. Open a command line prompt. Navigate to the *was_root/bin* directory. Type **wsadmin** to start the wsadmin command line tool.
2. Modify the eXtreme Scale PMI runtime configuration. Verify that PMI is enabled for the server using the following commands:

```
wsadmin>set s1 [$AdminConfig getid /Cell:CELL_NAME/Node:NODE_NAME/  
Server:APPLICATION_SERVER_NAME/  
wsadmin>set pmi [$AdminConfig list PMIService $s1]  
wsadmin>$AdminConfig show $pmi.
```

If PMI is not enabled, run the following commands to enable PMI:

```
wsadmin>$AdminConfig modify $pmi {{enable true}}  
wsadmin>$AdminConfig save
```

If you need to enable PMI, restart the server.

3. Set variables for changing the statistic set to a custom set using the following commands:

```
wsadmin>set perfName [$AdminControl completeObjectName type=Perf,  
process=APPLICATION_SERVER_NAME,*]  
wsadmin>set perfOName [$AdminControl makeObjectName $perfName]  
wsadmin>set params [java::new {java.lang.Object[]} 1]  
wsadmin>$params set 0 [java::new java.lang.String custom]  
wsadmin>set sigs [java::new {java.lang.String[]} 1]  
wsadmin>$sigs set 0 java.lang.String
```

4. Set statistic set to custom using the following command:

```
wsadmin>$AdminControl invoke_jmx $perfOName setStatisticSet $params $sigs
```

5. Set variables to enable the objectGridModule PMI statistic using the following commands:

```
wsadmin>set params [java::new {java.lang.Object[]} 2]  
wsadmin>$params set 0 [java::new java.lang.String objectGridModule=1]  
wsadmin>$params set 1 [java::new java.lang.Boolean false]  
wsadmin>set sigs [java::new {java.lang.String[]} 2]  
wsadmin>$sigs set 0 java.lang.String  
wsadmin>$sigs set 1 java.lang.Boolean
```

6. Set the statistics string using the following command:

```
wsadmin>set params2 [java::new {java.lang.Object[]} 2]
wsadmin>$params2 set 0 [java::new java.lang.String mapModule=*]
wsadmin>$params2 set 1 [java::new java.lang.Boolean false]
wsadmin>set sigs2 [java::new {java.lang.String[]} 2]
wsadmin>$sigs2 set 0 java.lang.String
wsadmin>$sigs2 set 1 java.lang.Boolean
```

7. Set the statistics string using the following command:

```
wsadmin>$AdminControl invoke_jmx $perf0Name setCustomSetString $params2 $sigs2
```

These steps enable eXtreme Scale runtime PMI, but do not modify the PMI configuration. If you restart the application server, the PMI settings are lost except for the main PMI enablement.

Example

You can perform the following steps to enable PMI statistics for the sample application:

1. Launch the application using the `http://host:port/ObjectGridSample` Web address, where `host` and `port` are the host name and HTTP port number of the server where the sample is installed.
2. In the sample application, click `ObjectGridCreationServlet`, and then click action buttons 1, 2, 3, 4, and 5 to generate actions to the ObjectGrid and maps. Do not close this servlet page right now.
3. In the administrative console, click **Monitoring and Tuning > Performance Monitoring Infrastructure > *server_name*** Click the **Runtime** tab.
4. Click the **Custom** radio button.
5. Expand the ObjectGrid Maps module in the runtime tree, then click the `clusterObjectGrid` link. Under the ObjectGrid Maps group, there is an ObjectGrid instance called `clusterObjectGrid`, and under the `clusterObjectGrid` group four maps exist: `counters`, `employees`, `offices`, and `sites`. In the ObjectGrids instance, there is the `clusterObjectGrid` instance, and under that instance is a transaction type called `DEFAULT`.
6. You can enable the statistics of interest to you. For example, you can enable number of map entries for `employees` map, and transaction response time for the `DEFAULT` transaction type.

What to do next

After PMI is enabled, you can view PMI statistics with the administrative console or through scripting.

Retrieving PMI statistics

By retrieving PMI statistics, you can see the performance of your eXtreme Scale applications.

Before you begin

- Enable PMI statistics tracking for your environment. See “Enabling PMI” on page 268 for more information.
- The paths in this task are assuming you are retrieving statistics for the sample application, but you can use these statistics for any other application with similar steps.

- If you are using the administrative console to retrieve statistics, you must be able to log in to the administrative console. If you are using scripting, you must be able to log in to wsadmin.

About this task

You can retrieve PMI statistics to view in Tivoli Performance Viewer by completing steps in the administrative console or with scripting.

- Administrative console steps
- Scripting steps

For more information about the statistics that can be retrieved, see “PMI modules” on page 272.

Procedure

- Retrieve PMI statistics in the administrative console.
 1. In the administrative console, click **Monitoring and tuning > Performance viewer > Current activity**
 2. Select the server that you want to monitor using Tivoli Performance Viewer, then enable the monitoring.
 3. Click the server to view the Performance viewer page.
 4. Expand the configuration tree. Click **ObjectGrid Maps > clusterObjectGrid** select **employees**. Expand **ObjectGrids > clusterObjectGrid** and select **DEFAULT**.
 5. In the ObjectGrid sample application, go to the ObjectGridCreationServlet servlet , click button 1, then populate maps. You can view the statistics in the viewer.
- Retrieve PMI statistics with scripting.
 1. On a command line prompt, navigate to the *was_root/bin* directory. Type `wsadmin` to start the wsadmin tool.
 2. Set variables for the environment using the following commands:


```
wsadmin>set perfName [$AdminControl completeObjectName type=Perf,*]
wsadmin>set perfOName [$AdminControl makeObjectName $perfName]
wsadmin>set mySrvName [$AdminControl completeObjectName type=Server,
name=APPLICATION_SERVER_NAME,*]
```
 3. Set variables to get mapModule statistics using the following commands:


```
wsadmin>set params [java::new {java.lang.Object[]} 3]
wsadmin>$params set 0 [$AdminControl makeObjectName $mySrvName]
wsadmin>$params set 1 [java::new java.lang.String mapModule]
wsadmin>$params set 2 [java::new java.lang.Boolean true]
wsadmin>set sigs [java::new {java.lang.String[]} 3]
wsadmin>$sigs set 0 javax.management.ObjectName
wsadmin>$sigs set 1 java.lang.String
wsadmin>$sigs set 2 java.lang.Boolean
```
 4. Get mapModule statistics using the following command:


```
wsadmin>$AdminControl invoke_jmx $perfOName getStatsString $params $sigs
```
 5. Set variables to get objectGridModule statistics using the following commands:


```
wsadmin>set params2 [java::new {java.lang.Object[]} 3]
wsadmin>$params2 set 0 [$AdminControl makeObjectName $mySrvName]
wsadmin>$params2 set 1 [java::new java.lang.String objectGridModule]
wsadmin>$params2 set 2 [java::new java.lang.Boolean true]
```

```

wsadmin>set sigs2 [java::new {java.lang.String[]} 3]
wsadmin>$sigs2 set 0 javax.management.ObjectName
wsadmin>$sigs2 set 1 java.lang.String
wsadmin>$sigs2 set 2 java.lang.Boolean

```

6. Get objectGridModule statistics using the following command:

```

wsadmin>$AdminControl invoke_jmx $perfObjectName getStatsString $params2 $sigs2

```

Results

You can view statistics in the Tivoli Performance Viewer.

PMI modules

You can monitor the performance of your applications with the performance monitoring infrastructure (PMI) modules.

objectGridModule

The objectGridModule contains a time statistic: transaction response time. A transaction is defined as the duration between the Session.begin method call and the Session.commit method call. This duration is tracked as the transaction response time. The root element of the objectGridModule, "root", serves as the entry point to the WebSphere eXtreme Scale statistics. This root element has ObjectGrids as its child elements, which have transaction types as their child elements. The response time statistic is associated with each transaction type.

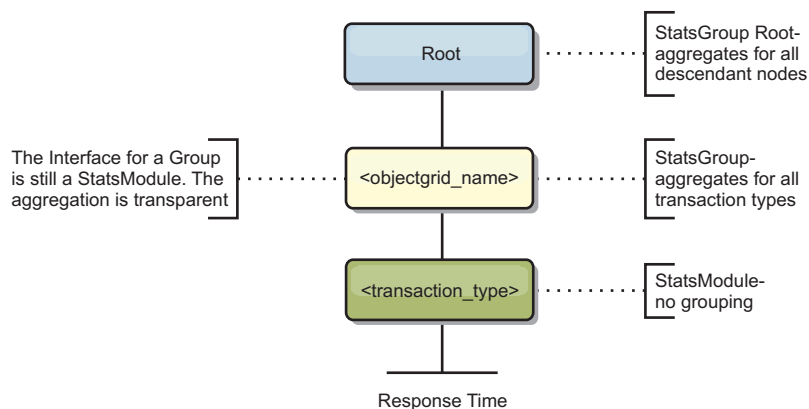


Figure 4. ObjectGridModule module structure

The following diagram shows an example of the ObjectGridModule structure. In this example, two ObjectGrid instances exist in the system: ObjectGrid A and ObjectGrid B. The ObjectGrid A instance has two types of transactions: A and default. The ObjectGrid B instance has only the default transaction type.

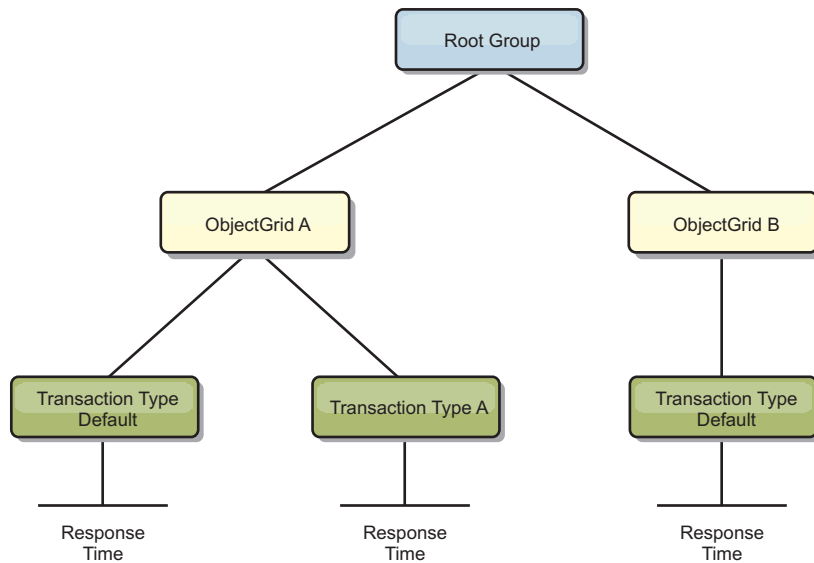


Figure 5. ObjectGridModule module structure example

Transaction types are defined by application developers because they know what types of transactions their applications use. The transaction type is set using the following `Session.setTransactionType(String)` method:

```

/**
 * Sets the transaction type for future transactions.
 *
 * After this method is called, all of the future transactions have the
 * same type until another transaction type is set. If no transaction
 * type is set, the default TRANSACTION_TYPE_DEFAULT transaction type
 * is used.
 *
 * Transaction types are used mainly for statistical data tracking purpose.
 * Users can predefine types of transactions that run in an
 * application. The idea is to categorize transactions with the same characteristics
 * to one category (type), so one transaction response time statistic can be
 * used to track each transaction type.
 *
 * This tracking is useful when your application has different types of
 * transactions.
 * Among them, some types of transactions, such as update transactions, process
 * longer than other transactions, such as read-only transactions. By using the
 * transaction type, different transactions are tracked by different statistics,
 * so the statistics can be more useful.
 *
 * @param tranType the transaction type for future transactions.
 */
void setTransactionType(String tranType);

```

The following example sets transaction type to `updatePrice`:

```

// Set the transaction type to updatePrice
// The time between session.begin() and session.commit() will be
// tracked in the time statistic for "updatePrice".
session.setTransactionType("updatePrice");
session.begin();
map.update(stockId, new Integer(100));
session.commit();

```

The first line indicates that the subsequent transaction type is `updatePrice`. An `updatePrice` statistic exists under the `ObjectGrid` instance that corresponds to the session in the example. Using Java Management Extensions (JMX) interfaces, you

can get the transaction response time for updatePrice transactions. You can also get the aggregated statistic for all types of transactions on the specified ObjectGrid instance.

mapModule

The mapModule contains three statistics that are related to eXtreme Scale maps:

- **Map hit rate** - *BoundedRangeStatistic*: Tracks the hit rate of a map. Hit rate is a float value between 0 and 100 inclusively, which represents the percentage of map hits in relation to map get operations.
- **Number of entries**-*CountStatistic*: Tracks the number of entries in the map.
- **Loader batch update response time**-*TimeStatistic*: Tracks the response time that is used for the loader batch-update operation.

The root element of the mapModule, "root", serves as the entry point to the ObjectGrid Map statistics. This root element has ObjectGrids as its child elements, which have maps as their child elements. Every map instance has the three listed statistics. The mapModule structure is shown in the following diagram:

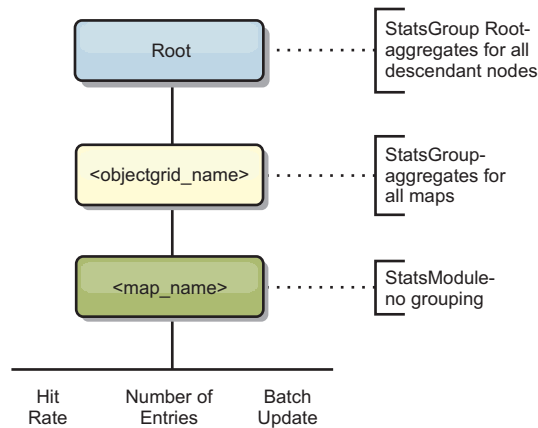
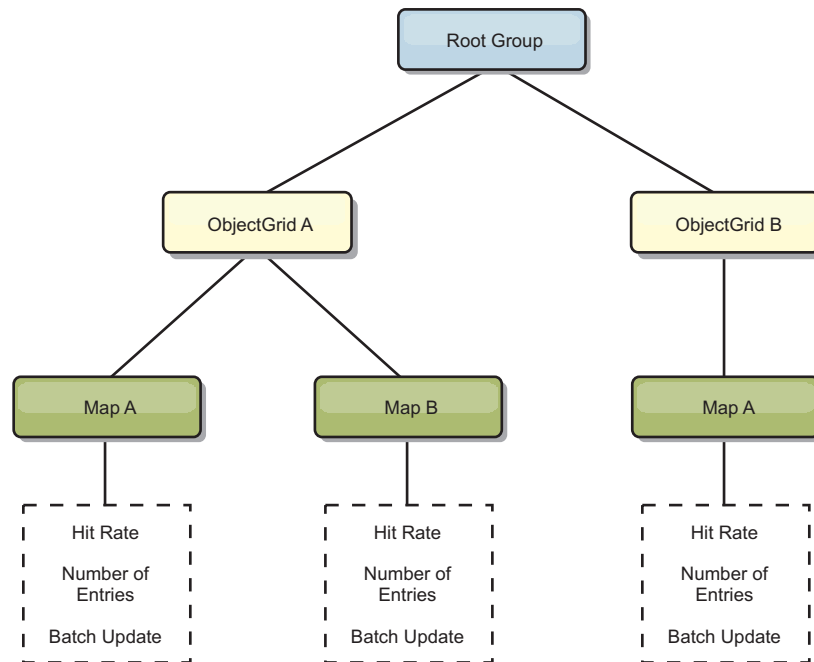


Figure 6. mapModule structure

The following diagram shows an example of the mapModule structure:

Figure 7. mapModule module structure example



hashIndexModule

The hashIndexModule contains the following statistics that are related to Map-level indexes:

- **Find Count-CountStatistic:** The number of invocations for the index find operation.
- **Collision Count-CountStatistic:** The number of collisions for the find operation.
- **Failure Count-CountStatistic:** The number of failures for the find operation.
- **Result Count-CountStatistic:** The number of keys returned from the find operation.
- **BatchUpdate Count-CountStatistic:** The number of batch updates against this index. When the corresponding map is changed in any manner, the index will have its doBatchUpdate() method called. This statistic will tell you how frequently your index is changing or being updated.
- **Find Operation Duration Time-TimeStatistic:** The amount of time the find operation takes to complete

The root element of the hashIndexModule, "root", serves as the entry point to the HashIndex statistics. This root element has ObjectGrids as its child elements, ObjectGrids have maps as their child elements, which finally have HashIndexes as their child elements and leaf nodes of the tree. Every HashIndex instance has the three listed statistics. The hashIndexModule structure is shown in the following diagram:

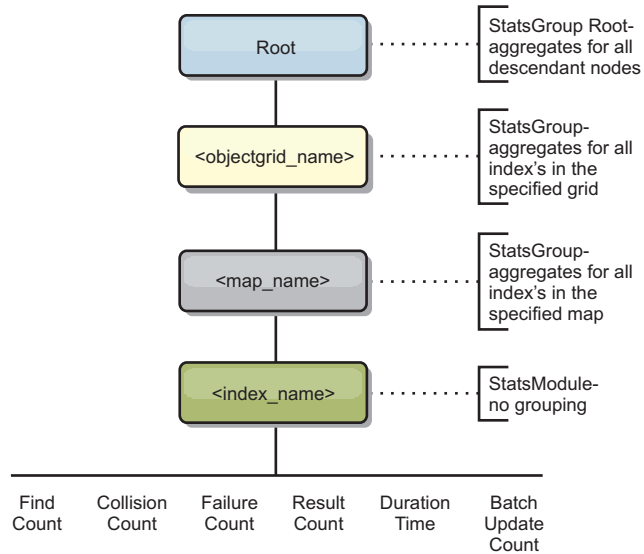


Figure 8. hashIndexModule module structure

The following diagram shows an example of the hashIndexModule structure:

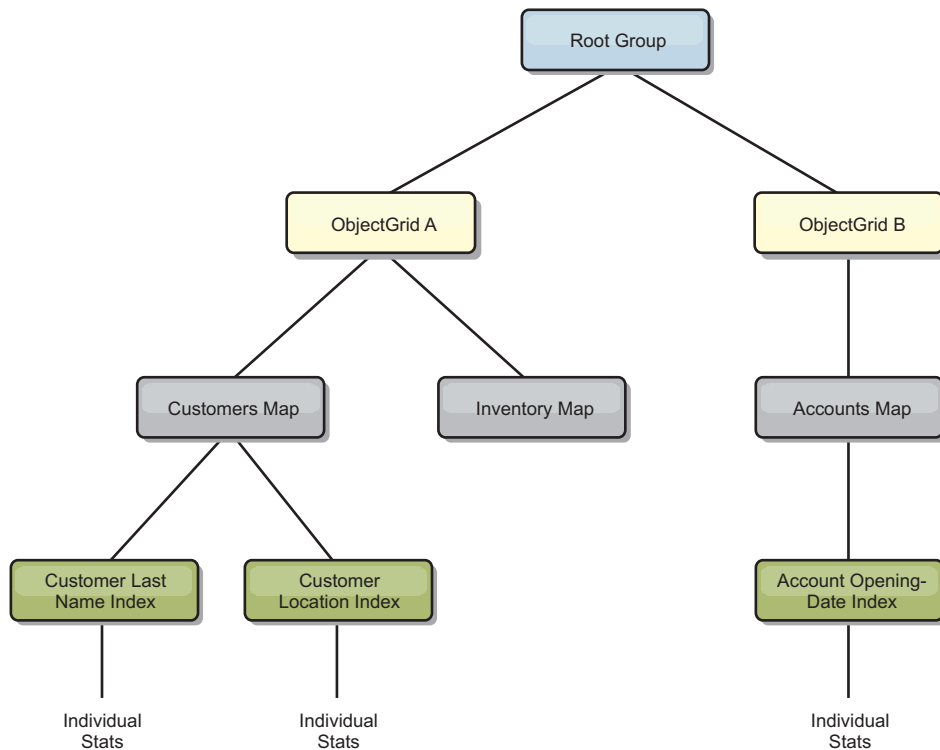


Figure 9. hashIndexModule module structure example

agentManagerModule

The agentManagerModule contains statistics that are related to map-level agents:

- **Reduce Time:** *TimeStatistic* - The amount of time for the agent to finish the reduce operation.
- **Total Duration Time:** *TimeStatistic* - The total amount of time for the agent to complete all operations.

- **Agent Serialization Time:** *TimeStatistic* - The amount of time to serialize the agent.
- **Agent Inflation Time:** *TimeStatistic* - The amount of time it takes to inflate the agent on the server.
- **Result Serialization Time:** *TimeStatistic* - The amount of time to serialize the results from the agent.
- **Result Inflation Time:** *TimeStatistic* - The amount of time to inflate the results from the agent.
- **Failure Count:** *CountStatistic* - The number of times that the agent failed.
- **Invocation Count:** *CountStatistic* - The number of times the AgentManager has been invoked.
- **Partition Count:** *CountStatistic* - The number of partitions to which the agent is sent.

The root element of the agentManagerModule, "root", serves as the entry point to the AgentManager statistics. This root element has ObjectGrids as its child elements, ObjectGrids have maps as their child elements, which finally have AgentManager instances as their child elements and leaf nodes of the tree. Every AgentManager instance has statistics.

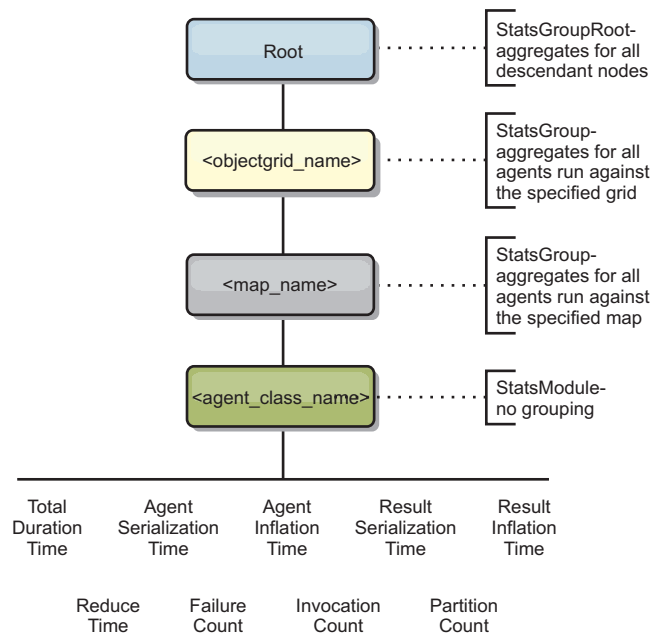


Figure 10. agentManagerModule structure

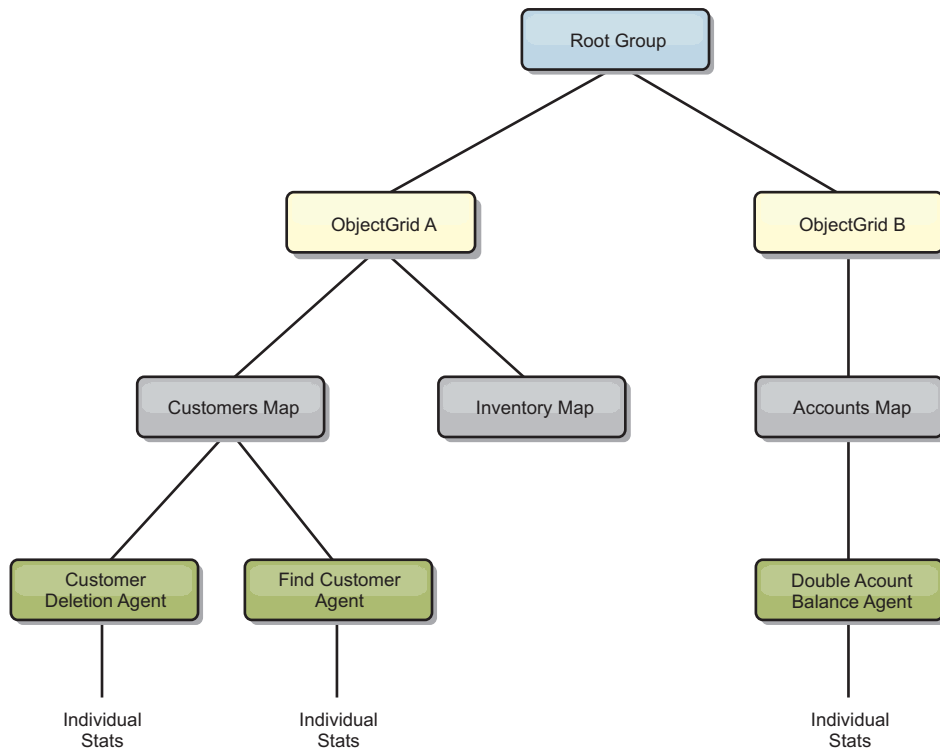


Figure 11. agentManagerModule structure example

queryModule

The queryModule contains statistics that are related to eXtreme Scale queries:

- **Plan Creation Time:** *TimeStatistic* - The amount of time to create the query plan.
- **Execution Time:** *TimeStatistic* - The amount of time to run the query.
- **Execution Count:** *CountStatistic* - The number of times the query has been run.
- **Result Count:** *CountStatistic* - The count for each the result set of each query run.
- **FailureCount:** *CountStatistic* - The number of times the query has failed.

The root element of the queryModule, "root", serves as the entry point to the Query Statistics. This root element has ObjectGrids as its child elements, which have Query objects as their child elements and leaf nodes of the tree. Every Query instance has the three listed statistics.

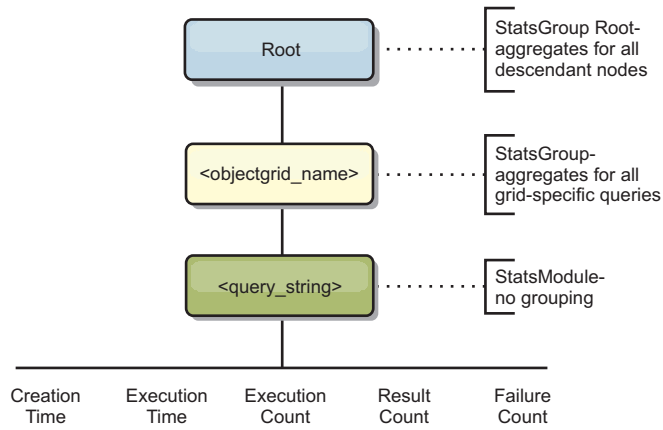


Figure 12. queryModule structure

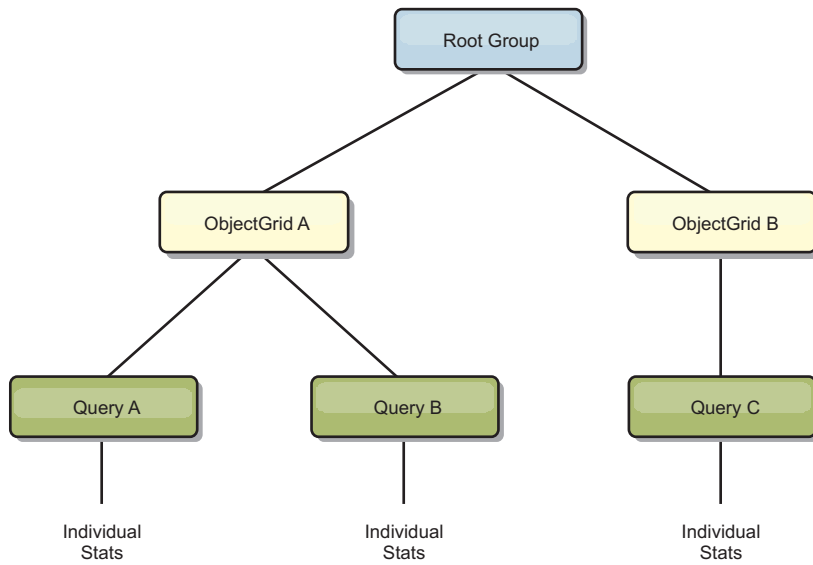


Figure 13. QueryStats.jpg queryModule structure example

Chapter 6. Programming for JPA integration

The Java Persistence API (JPA) is a specification that allows mapping Java objects to relational databases. JPA contains a full object-relational mapping (ORM) specification using Java language metadata annotations, XML descriptors, or both to define the mapping between Java objects and a relational database. A number of open-source and commercial implementations are available.

To use JPA, you must have a supported JPA provider, such as OpenJPA or Hibernate, JAR files, and a META-INF/persistence.xml file in your class path.

JPA Loaders

The Java Persistence API (JPA) is a specification that allows mapping Java objects to relational databases. JPA contains a full object-relational mapping (ORM) specification using Java language metadata annotations, XML descriptors, or both to define the mapping between Java objects and a relational database. A number of open-source and commercial implementations are available.

You can use a Java Persistence API (JPA) loader plug-in implementation with eXtreme Scale to interact with any database supported by your chosen loader. To use JPA, you must have a supported JPA provider, such as OpenJPA or Hibernate, JAR files, and a META-INF/persistence.xml file in your class path.

The JPALoader `com.ibm.websphere.objectgrid.jpa.JPALoader` and the JPAEntityLoader `com.ibm.websphere.objectgrid.jpa.JPAEntityLoader` plug-ins are two built-in JPA loader plug-ins that are used to synchronize the ObjectGrid maps with a database. You must have a JPA implementation, such as Hibernate or OpenJPA, to use this feature. The database can be any back end that is supported by the chosen JPA provider.

You can use the JPALoader plug-in when you are storing data using the ObjectMap API. Use the JPAEntityLoader plug-in when you are storing data using the EntityManager API.

JPA loader architecture

The JPA Loader is used for eXtreme Scale maps that store plain old Java objects (POJO).




Figure 14. JPA Loader architecture

When an `ObjectMap.get(Object key)` method is called, the eXtreme Scale run time first checks whether the entry is contained in the `ObjectMap` layer. If not, the run time delegates the request to the JPA Loader. Upon request of loading the key, the `JPALoader` calls the `JPA EntityManager.find(Object key)` method to find the data from the JPA layer. If the data is contained in the JPA entity manager, it is returned; otherwise, the JPA provider interacts with the database to get the value.

When an update to `ObjectMap` occurs, for example, using the `ObjectMap.update(Object key, Object value)` method, the eXtreme Scale run time creates a `LogElement` for this update and sends it to the `JPALoader`. The `JPALoader` calls the `JPA EntityManager.merge(Object value)` method to update the value to the database.

For the `JPAEntityLoader`, the same four layers are involved. However, because the `JPAEntityLoader` plug-in is used for maps that store eXtreme Scale entities, relations among entities could complicate the usage scenario. An eXtreme Scale entity is distinguished from JPA entity. For more details, see “`JPAEntityLoader` plug-in” on page 238.

Methods

Loaders provide three main methods:

1. `get`: Returns a list of values that correspond to the list of keys that are passed in by retrieving the data using JPA. The method uses JPA to find the entities in the database. For the `JPALoader` plug-in, the returned list contains a list of JPA entities directly from the find operation. For the `JPAEntityLoader` plug-in, the returned list contains eXtreme Scale entity value tuples that are converted from the JPA entities.
2. `batchUpdate`: Writes the data from `ObjectGrid` maps to the database. Depending on different operation types (insert, update, or delete), the loader

uses the JPA persist, merge, and remove operations to update the data to the database. For the JPALoader, the objects in the map are directly used as JPA entities. For the JPAEntityLoader, the entity tuples in the map are converted into objects which are used as JPA entities.

3. `preloadMap`: Preloads the map using the `ClientLoader.load` client loader method. For partitioned maps, the `preloadMap` method is only called in one partition. The partition is specified the `preloadPartition` property of the JPALoader or JPAEntityLoader class. If the `preloadPartition` value is set to less than zero, or greater than $(total_number_of_partitions - 1)$, preload is disabled.

Both JPALoader and JPAEntityLoader plug-ins work with the JPATxCallback class to coordinate the eXtreme Scale transactions and JPA transactions. JPATxCallback must be configured in the ObjectGrid instance to use these two loaders.

Developing client-based JPA loaders

You can implement preloading and reloading of data in your application with a Java Persistence API (JPA) utility. This capability can simplify loading the maps when the queries to the database cannot be partitioned.

Before you begin

- You must be using a JPA provider with a supported database.
- Before you preload or reload maps, you must set the availability state of the ObjectGrid to PRELOAD. You can set the availability state with the `setObjectGridState` method of the StateManager interface. The StateManager interface prevents other clients from accessing the ObjectGrid when it is not yet online. After you preload or reload the map, you can set the state back to ONLINE.
- When you are preloading different maps in one ObjectGrid, set the ObjectGrid state to PRELOAD one time and set the value back to ONLINE after all maps finish data loading. This coordination can be done by the ClientLoadCallback interface. Set the ObjectGrid state to PRELOAD after the first `preStart` notification from the ObjectGrid instance, and set it back to ONLINE after the last `postFinish` notification.
- If you need to preload maps from different Java virtual machines, you have to coordinate between multiple Java virtual machines. Set the ObjectGrid state to PRELOAD one time before the first map is being preloaded in any of the Java virtual machines, and set the value back to ONLINE after all maps finish data loading in all the Java virtual machines.

the information about setting the availability of an ObjectGrid in the *Administration Guide* for more information

About this task

When you run a preload or reload operation on your map, the following actions occur:

1. The initial action that is taken depends on if you are running a preload or reload operation.
 - **Preload operation:** The map to be preloaded is cleared. For an entity map, if any relation is configured as cascade-remove, any related maps are cleared.

- **Reload operation:** The provided query is run on the map and the results are invalidated. For an entity map, if any relation is configured with the **CascadeType.INVALIDATE** option, the related entities are also invalidated from their maps.
2. Run the query to JPA for the entities in a batch.
 3. For each batch, a key list and value list for each partition is built.
 4. For each partition, the data grid agent is called to insert or update the data on the server side directly if it is an eXtreme Scale client. If the data grid is a local instance, the data in the maps is directly inserted or updated.

Client-based JPA preload utility overview

The client-based Java Persistence API (JPA) preload utility loads data into eXtreme Scale backing maps using a client connection to the ObjectGrid.

This capability can simplify loading the maps when the queries to the database cannot be partitioned. A loader, such as a JPA Loader can also be used and is ideal when the data can be loaded in parallel.

The client-based JPA preload utility can use either the OpenJPA or Hibernate JPA implementations to load the ObjectGrid from a database. Because WebSphere eXtreme Scale does not directly interact with the database or Java Database Connectivity (JDBC), any database that OpenJPA or Hibernate supports can be used to load the ObjectGrid.




Figure 15. Client loader that uses JPA implementation to load the ObjectGrid

Typically, a user application provides a persistence unit name, an entity class name, and a JPA query to the client loader. The client loader retrieves the JPA entity manager based on the persistence unit name, uses the entity manager to query data from the database with the provided entity class and JPA query, and finally loads the data into the distributed ObjectGrid maps. When multi-level relations are

involved in the query, can use a custom query string to optimize the performance. Optionally, a persistence property map could be provided to override the configured persistence properties.

A client loader can load data in two different modes, as displayed in the following table:

Table 5. Client loader modes

| Mode | Description |
|----------------|--|
| <i>Preload</i> | Clears and loads all entries into the backing map. If the map is an entity map, any related entity maps will also be cleared if the ObjectGrid CascadeType.REMOVE option is enabled. |
| <i>Reload</i> | The JPA query is executed against the ObjectGrid to invalidate all the entities in the map that match the query. If the map is an entity map, any related entity maps will also be cleared if the ObjectGrid CascadeType.INVALIDATE option is enabled. |

In either case, a JPA query is used to select and load the desired entities from the database and to store them in the ObjectGrid maps. If the ObjectGrid map is a non-entity map, the JPA entities will be detached and stored directly. If the ObjectGrid map is an entity map, the JPA entities are stored as ObjectGrid entity tuples. You can provide a JPA query or use the default query `select o from EntityName o`.

For more information about configuring the client-based JPA preload utility, see the information in the *Programming Guide*

Example: Preloading a map with the ClientLoader interface

You can preload a map to populate the map data before clients begin accessing the map.

Client-based preload example

The following sample code snippet shows a simple client loading. In this example, the CUSTOMER map is configured as an entity map. The Customer entity class, which is configured in the ObjectGrid entity metadata descriptor XML file, has a one-to-many relation with Order entities. The Customer entity has the CascadeType.ALL option enabled on the relation to the Order entity. Before the ClientLoader.load method is called, the ObjectGrid state is set to PRELOAD. The **isPreload** parameter on the load method is set to true.

```
// Get the StateManager
StateManager stateMgr = StateManagerFactory.getStateManager();

// Set ObjectGrid state to PRELOAD before calling ClientLoader.load
stateMgr.setObjectGridState(AvailabilityState.PRELOAD, objectGrid);

ClientLoader c = ClientLoaderFactory.getClientLoader();

// Load the data
c.load(objectGrid, "CUSTOMER", "customerPU", null,
    null, null, null, true, null);
```

```
// Set ObjectGrid state back to ONLINE
stateMgr.setObjectGridState(AvailabilityState.ONLINE, objectGrid);
```

Example: Reloading a map with the ClientLoader interface

Reloading a map is the same as preloading a map, except that the `isPreload` argument is set to `false` in the `ClientLoader.load` method.

Client-based reload example

The following sample shows how to reload maps. Compared to the preload sample, the main difference is that a `loadSql` and parameters are provided. This sample only reloads the Customer data with an ID between 1000 and 2000. The `isPreload` parameter on the load method is set to `false`.

```
// Get the StateManager
StateManager stateMgr = StateManagerFactory.getStateManager();

// Set ObjectGrid state to PRELOAD before calling ClientLoader.loader
stateMgr.setObjectGridState(AvailabilityState.PRELOAD, objectGrid);

ClientLoader c = ClientLoaderFactory.getClientLoader();

// Load the data
String loadSql = "select c from CUSTOMER c
  where c.custId >= :startCustId and c.custId < :endCustId ";
Map<String, Long> params = new HashMap<String, Long>();
params.put("startCustId", 1000L);
params.put("endCustId", 2000L);

c.load(objectGrid, "CUSTOMER", "customerPU", null, null,
  loadSql, params, false, null);

// Set ObjectGrid state back to ONLINE
stateMgr.setObjectGridState(AvailabilityState.ONLINE, objectGrid);
```

Remember: This query string observes both JPA query syntax and eXtreme Scale entity query syntax. This query string is important because it runs twice: to invalidate the matched ObjectGrid entities and to load the matched JPA entities.

Example: Calling a client loader

You can use the `preload` method in the `Loader` interface to call a client loader.

Use the `preload` method in the `Loader` interface to call a client loader:

```
void preloadMap(Session session, BackingMap backingMap) throws LoaderException;
```

This method signals the loader to preload the data into the map. A loader implementation can use a client loader to preload the data to all its partitions. For example, the JPA loader uses the client loader to preload data into the map.

For more information, see the JPA loaders overview topic in the *Product Overview*.

Example: Calling a client loader with the preloadMap method

An example of how to preload the map using the client loader in the `preloadMap` method follows. The example first checks whether the current partition number is the same as the preload partition. If the partition number is not the same as the

preload partition, no action occurs. If the partition numbers match, the client loader is called to load data into the maps. You must call the client loader in one and only one partition.

```
void preloadMap (Session session, BackingMap backingMap) throws LoaderException {
```

```
....
ObjectGrid objectGrid = session.getObjectGrid();
int partitionId = backingMap.getPartitionId();
int numPartitions = backingMap.getPartitionManager().getNumOfPartitions();

// Only call client loader data in one partition
if (partitionId == preloadPartition) {
    ClientLoader c = ClientLoaderFactory.getClientLoader();
    // Call the client loader to load the data
    try {
        c.load(objectGrid, "CUSTOMER", "customerPU",
            null, entityClass, null, null, true, null);
    } catch (ObjectGridException e) {
        LoaderException le = new LoaderException("Exception caught in ObjectMap " +
            ogName + "." + mapName);
        le.initCause(e);
        throw le;
    }
}
```

Remember: Configure the backingMap attribute "preloadMode" to true, so the preload method runs asynchronously. Otherwise, the preload method blocks the ObjectGrid instance from being activated.

Example: Creating a custom client-based JPA loader

The ClientLoader.load method in the Loader interface provides a client load function that satisfies most scenarios. However, if you want to load data without the ClientLoader.load method, you can implement your own preload utility.

Custom loader template

Use the following template to develop your loader:

```
// Get the StateManager
StateManager stateMgr = StateManagerFactory.getStateManager();

// Set ObjectGrid state to PRELOAD before calling ClientLoader.loader
stateMgr.setObjectGridState(AvailabilityState.PRELOAD, objectGrid);

// Load the data
...<your preload utility implementation>...

// Set ObjectGrid state back to ONLINE
stateMgr.setObjectGridState(AvailabilityState.ONLINE, objectGrid);
```

Developing a client-based JPA loader with a DataGrid agent

When loading from the client side, using a DataGrid agent could increase performance. By using a DataGrid agent, all the data reads and writes occur in the server process. You can also design your application to make sure that DataGrid agents on multiple partitions run in parallel to further boost performance.

About this task

For more information about the DataGrid agent, see “DataGrid APIs and partitioning” on page 155.

After you create the data preload implementation, you can create a generic Loader to complete the following tasks:

- Query the data from database in batches.
- Build a key list and value list for each partition.
- For each partition, call the `agentMgr.callReduceAgent(agent, aKey)` method to run the agent in the server in a thread. By running in a thread, you can run agents concurrently on multiple partitions.

Example

The following snippet of code is an example of how to load using a DataGrid agent:

```
import java.io.Externalizable;
import java.io.IOException;
import java.io.ObjectInput;
import java.io.ObjectOutput;
import java.util.ArrayList;
import java.util.Collection;
import java.util.Iterator;
import java.util.List;

import com.ibm.websphere.objectgrid.NoActiveTransactionException;
import com.ibm.websphere.objectgrid.ObjectGridException;
import com.ibm.websphere.objectgrid.ObjectGridRuntimeException;
import com.ibm.websphere.objectgrid.ObjectMap;
import com.ibm.websphere.objectgrid.Session;
import com.ibm.websphere.objectgrid.TransactionException;
import com.ibm.websphere.objectgrid.datagrid.ReduceGridAgent;
import com.ibm.websphere.objectgrid.em.EntityManager;

public class InsertAgent implements ReduceGridAgent, Externalizable {

    private static final long serialVersionUID = 6568906743945108310L;

    private List keys = null;

    private List vals = null;

    protected boolean isEntityMap;

    public InsertAgent() {
    }

    public InsertAgent(boolean entityMap) {
        isEntityMap = entityMap;
    }

    public Object reduce(Session sess, ObjectMap map) {
        throw new UnsupportedOperationException(
            "ReduceGridAgent.reduce(Session, ObjectMap)");
    }

    public Object reduce(Session sess, ObjectMap map, Collection arg2) {
        Session s = null;
        try {
            s = sess.getObjectGrid().getSession();
            ObjectMap m = s.getMap(map.getName());
            s.beginNoWriteThrough();
        }
    }
}
```

```

        Object ret = process(s, m);
        s.commit();
        return ret;
    } catch (ObjectGridRuntimeException e) {
        if (s.isTransactionActive()) {
            try {
                s.rollback();
            } catch (TransactionException e1) {
            } catch (NoActiveTransactionException e1) {
            }
        }
        throw e;
    } catch (Throwable t) {
        if (s.isTransactionActive()) {
            try {
                s.rollback();
            } catch (TransactionException e1) {
            } catch (NoActiveTransactionException e1) {
            }
        }
        throw new ObjectGridRuntimeException(t);
    }
}

public Object process(Session s, ObjectMap m) {
    try {

        if (!isEntityMap) {
            // In the POJO case, it is very straightforward,
            // we can just put everything in the
            // map using insert
            insert(m);
        } else {
            // 2. Entity case.
            // In the Entity case, we can persist the entities
            EntityManager em = s.getEntityManager();
            persistEntities(em);
        }
        return Boolean.TRUE;
    } catch (ObjectGridRuntimeException e) {
        throw e;
    } catch (ObjectGridException e) {
        throw new ObjectGridRuntimeException(e);
    } catch (Throwable t) {
        throw new ObjectGridRuntimeException(t);
    }
}

/**
 * Basically this is fresh load.
 * @param s
 * @param m
 * @throws ObjectGridException
 */
protected void insert(ObjectMap m) throws ObjectGridException {

    int size = keys.size();

    for (int i = 0; i < size; i++) {
        m.insert(keys.get(i), vals.get(i));
    }
}
}

```

```

protected void persistEntities(EntityManager em) {
    Iterator<Object> iter = vals.iterator();

    while (iter.hasNext()) {
        Object value = iter.next();
        em.persist(value);
    }
}

public Object reduceResults(Collection arg0) {
    return arg0;
}

public void readExternal(ObjectInput in)
    throws IOException, ClassNotFoundException {
    int v = in.readByte();
    isEntityMap = in.readBoolean();
    vals = readList(in);
    if (!isEntityMap) {
        keys = readList(in);
    }
}

public void writeExternal(ObjectOutput out) throws IOException {
    out.write(1);
    out.writeBoolean(isEntityMap);

    writeList(out, vals);
    if (!isEntityMap) {
        writeList(out, keys);
    }
}

public void setData(List ks, List vs) {
    vals = vs;
    if (!isEntityMap) {
        keys = ks;
    }
}

/**
 * @return Returns the isEntityMap.
 */
public boolean isEntityMap() {
    return isEntityMap;
}

static public void writeList(ObjectOutput oo, Collection l)
    throws IOException {
    int size = l == null ? -1 : l.size();
    oo.writeInt(size);
    if (size > 0) {
        Iterator iter = l.iterator();
        while (iter.hasNext()) {
            Object o = iter.next();
            oo.writeObject(o);
        }
    }
}

public static List readList(ObjectInput oi)
    throws IOException, ClassNotFoundException {
    int size = oi.readInt();
    if (size == -1) {
        return null;
    }
}

```

```

    }

    ArrayList list = new ArrayList(size);
    for (int i = 0; i < size; ++i) {
        Object o = oi.readObject();
        list.add(o);
    }
    return list;
}
}

```

Starting the JPA time-based updater

When you start the Java Persistence API (JPA) time-based updater, the ObjectGrid maps are updated with the latest changes in the database.

Before you begin

Configure the time-based updater. See the information about configuring a JPA time-based data updater in the *Administration Guide*.

About this task

For more information about how the Java Persistence API (JPA) time-based data updater works, see “JPA time-based data updater” on page 294.

Procedure

- Start a time-based database updater.
 - **Automatically for distributed eXtreme Scale:**
If you create the timeBasedDBUpdate configuration for the backing map, the time-based database updater is automatically started when a distributed ObjectGrid primary shard is activated. For an ObjectGrid that has multiple partitions, the time-based database updater only starts at partition 0.
 - **Automatically for local eXtreme Scale:**
If you create the timeBasedDBUpdate configuration for the backing map, the time-based database updater is automatically started when the local map is activated.
 - **Manually:**
You can also manually start or stop a time-based database updater using the TimeBasedDBUpdater API.


```
public synchronized void startDBUpdate(ObjectGrid objectGrid, String mapName,
String punitName, Class entityClass, String timestampField, DBUpdateMode mode) {
```

 1. **ObjectGrid:** the ObjectGrid instance (local or client).
 2. **mapName:** the name of the backing map for which the time-based database updater is started.
 3. **punitName:** the JPA persistence unit name for creating a JPA entity manager factory; the default value is the name of the first persistence unit defined in the persistence.xml file.
 4. **entityClass:** The entity class name used to interact with the Java Persistence API (JPA) provider; the entity class name is used to get JPA entities using entity queries.
 5. **timestampField:** A timestamp field of the entity class to identify the time or sequence when a database back end record was last updated or inserted.

6. **mode:** The time-based database update mode; an `INVALIDATE_ONLY` type causes it to invalidate the entries in the ObjectGrid map if the corresponding records in the database have changed; an `UPDATE_ONLY` type indicates to update the existing entries in the ObjectGrid map with the latest values from the database; however, all the newly inserted records to the database are ignored; an `INSERT_UPDATE` type indicates to update the existing entries in the ObjectGrid map with the latest values from the database; also, all the newly inserted records to the database are inserted into the ObjectGrid map.

If you want to stop the time-based database updater, you can call the following method to stop the updater:

```
public synchronized void stopDBUpdate(ObjectGrid objectGrid, String mapName)
```

The ObjectGrid and mapName parameters should be the same as those passed in the startDBUpdate method.

- Create the timestamp field in your database.

– DB2

As a part of the optimistic locking feature, DB2 9.5 provides a row change timestamp feature. You can define a column ROWCHGTS using the ROW CHANGE TIMESTAMP format as follows:

```
ROWCHGTS TIMESTAMP NOT NULL
GENERATED ALWAYS
FOR EACH ROW ON UPDATE AS
ROW CHANGE TIMESTAMP
```

Then you can indicate the entity field which corresponds to this column as the timestamp field by either annotation or configuration. An example follows:

```
@Entity(name = "USER_DB2")
@Table(name = "USER1")
public class User_DB2 implements Serializable {

    private static final long serialVersionUID = 1L;

    public User_DB2() {
    }

    public User_DB2(int id, String firstName, String lastName) {
        this.id = id;
        this.firstName = firstName;
        this.lastName = lastName;
    }

    @Id
    @Column(name = "ID")
    public int id;

    @Column(name = "FIRSTNAME")
    public String firstName;

    @Column(name = "LASTNAME")
    public String lastName;

    @com.ibm.websphere.objectgrid.jpa.dbupdate.annotation.Timestamp
    @Column(name = "ROWCHGTS", updatable = false, insertable = false)
    public Timestamp rowChgTs;
}
```

– Oracle

In Oracle, there is a pseudo-column `ora_rowscn` for the system change number of the record. You can use this column for the same purpose. An

example of the entity that uses the ora_rowscn field as the time-based database update timestamp field follows:

```
@Entity(name = "USER_ORA")
@Table(name = "USER1")
public class User_ORA implements Serializable {

    private static final long serialVersionUID = 1L;

    public User_ORA() {
    }

    public User_ORA(int id, String firstName, String lastName) {
        this.id = id;
        this.firstName = firstName;
        this.lastName = lastName;
    }

    @Id
    @Column(name = "ID")
    public int id;

    @Column(name = "FIRSTNAME")
    public String firstName;

    @Column(name = "LASTNAME")
    public String lastName;

    @com.ibm.websphere.objectgrid.jpa.dbupdate.annotation.Timestamp
    @Column(name = "ora_rowscn", updatable = false, insertable = false)
    public long rowChgTs;
}
```

– Other databases

For other types of databases, you can create a table column to track the changes. The column values have to be manually managed by the application that updates the table.

Take an Apache Derby database as an example: You can create a column "ROWCHGTS" to track the change numbers. Also, a latest change number is tracked for this table. Every time a record is inserted or updated, the latest change number for the table is incremented, and the ROWCHGTS column value for the record is updated with this incremented number.

```
@Entity(name = "USER_DER")
@Table(name = "USER1")
public class User_DER implements Serializable {

    private static final long serialVersionUID = 1L;

    public User_DER() {
    }

    public User_DER(int id, String firstName, String lastName) {
        this.id = id;
        this.firstName = firstName;
        this.lastName = lastName;
    }

    @Id
    @Column(name = "ID")
    public int id;

    @Column(name = "FIRSTNAME")
    public String firstName;

    @Column(name = "LASTNAME")
```

```

public String lastName;

@com.ibm.websphere.objectgrid.jpa.dbupdate.annotation.Timestamp
@Column(name = "ROWCHGTS", updatable = true, insertable = true)
public long rowChgTs;
}

```

JPA time-based data updater

A Java Persistence API (JPA) time-based database updater updates the ObjectGrid maps with the latest changes in the database.

When changes are made directly to a database that is being fronted by WebSphere eXtreme Scale, those changes are not concurrently reflected in the eXtreme Scale grid. To properly implement eXtreme Scale as an in-memory database processing space, take into consideration that your grid can get out of sync with the database. Time-based database updater uses the System Change Number (SCN) capability in Oracle 10g and row change timestamp in DB2 9.5 to monitor changes in the database for invalidation and update. The updater also allows applications to have a user-defined field for the same purpose.




Figure 16. Periodic refresh

The time-based database updater periodically queries the database using JPA interfaces to get the JPA entities that represent the newly inserted and updated records in the database. To periodically update the records, every record in the database should have a timestamp to identify the time or sequence in which the record was last updated or inserted. It is not required that the timestamp be in a timestamp format. The timestamp value can be in an integer or long format, if it generates a unique, increasing value.

Several commercial databases have provided this capability.

For example, in DB2 9.5, you can define a column using the ROW CHANGE TIMESTAMP format as follows:

```

ROWCHGTS TIMESTAMP NOT NULL
GENERATED ALWAYS
FOR EACH ROW ON UPDATE AS
ROW CHANGE TIMESTAMP

```

In Oracle, you can use the pseudo-column **ora_rowscn**, which represents the system change number of the record.

The time-based database updater updates the ObjectGrid maps in three different ways:

1. `INVALIDATE_ONLY`. Invalidate the entries in the ObjectGrid map if the corresponding records in the database have changed.
2. `UPDATE_ONLY`. Update the entries in the ObjectGrid map if the corresponding records in the database have changed. However, all the newly inserted records to the database are ignored.
3. `INSERT_UPDATE`. Update the existing entries in the ObjectGrid map with the latest values from the database. Also, all the newly inserted records to the database are inserted into the ObjectGrid map.

For more information about configuring the JPA time-based data updater, see the information in the *Administration Guide*.

Chapter 7. Programming for Spring integration

Learn how to integrate your eXtreme Scale applications with the popular Spring framework.

Spring framework

Spring is a framework for developing Java applications. WebSphere eXtreme Scale provides support to allow Spring to manage transactions and configure the clients and servers comprising your deployed in-memory data grid.

Spring managed native transactions

Spring provides container-managed transactions that are similar to a Java Platform, Enterprise Edition application server. However, the Spring mechanism can use different implementations. WebSphere eXtreme Scale provides transaction manager integration which allows Spring to manage the ObjectGrid transaction life cycles. See the information about native transactions in the *Programming Guide* for details.

Spring managed extension beans and namespace support

Also, eXtreme Scale integrates with Spring to allow Spring-style beans defined for extension points or plug-ins. This feature provides more sophisticated configurations and more flexibility for configuring the extension points.

In addition to Spring managed extension beans, eXtreme Scale provides a Spring namespace called "objectgrid". Beans and built-in implementations are pre-defined in this namespace, which makes it easier for users to configure eXtreme Scale.

Shard scope support

With the traditional style Spring configuration, an ObjectGrid bean can either be a singleton type or prototype type. ObjectGrid also supports a new scope called the "shard" scope. If a bean is defined as shard scope, then only one bean is created per shard. All requests for beans with an ID or IDs matching that bean definition in the same shard results in that one specific bean instance being returned by the Spring container.

The following example shows that a `com.ibm.ws.objectgrid.jpa.plugins.JPAPropFactoryImpl` bean is defined with scope set to `shard`. Therefore, only one instance of the `JPAPropFactoryImpl` class is created per shard.

```
<bean id="jpaPropFactory" class="com.ibm.ws.objectgrid.jpa.plugins.JPAPropFactoryImpl" scope="shard" />
```

Spring Web Flow

Spring Web Flow stores its session state in an HTTP session by default. If a web application uses eXtreme Scale for session management, then Spring automatically stores state with eXtreme Scale. Also, fault tolerance is enabled in the same manner as the session.

See the ObjectMap API information in the *Product Overview* for further details.

Packaging

The eXtreme Scale Spring extensions are in the ogspring.jar file. This Java archive (JAR) file must be on the class path for Spring support to work. If a Java EE application that is running in a WebSphere Extended Deployment augmented WebSphere Application Server Network Deployment, put the spring.jar file and its associated files in the enterprise archive (EAR) modules. You must also place the ogspring.jar file in the same location.

Managing transactions with Spring

Spring is a popular framework for developing Java applications. WebSphere eXtreme Scale provides support to allow Spring to manage eXtreme Scale transactions and configure eXtreme Scale clients and servers.

About this task

The Spring Framework is highly integrable with eXtreme Scale, as discussed in the following sections.

Procedure

- **Native transactions:** Spring provides container-managed transactions along the style of a Java Platform, Enterprise Edition application server but has the advantage that Springs mechanism can have different implementations plugged in. This topic describes an eXtreme Scale Platform Transaction manager that can be used with Spring. This allows programmers to annotate their POJOs (plain old Java objects) and then have Spring automatically acquire Sessions from eXtreme Scale and begin, commit, rollback, suspend, and resume eXtreme Scale transactions. Spring transactions are described more fully in Chapter 10 of the official Spring reference documentation. The following explains how to create an eXtreme Scale transaction manager and use it with annotated POJOs. It also explains how to use this approach with client or local eXtreme Scale as well as a collocated Data Grid style application.
- **Transaction manager:** To work with Spring,, eXtreme Scale provides an implementation of a Spring PlatformTransactionManager. This manager can provide managed eXtreme Scale sessions to POJOs managed by Spring. Through the use of annotations, Spring manages those sessions for the POJOs in terms of transaction life cycle. The following XML snippet shows how to create a transaction Manager:

```
<aop:aspectj-autoproxy/>
<tx:annotation-driven transaction-manager="transactionManager"/>

<bean id="ObjectGridManager"
      class="com.ibm.websphere.objectgrid.ObjectGridManagerFactory"
      factory-method="getObjectGridManager"/>

<bean id="ObjectGrid"
      factory-bean="ObjectGridManager"
      factory-method="createObjectGrid"/>

<bean id="transactionManager"
      class="com.ibm.websphere.objectgrid.spring.ObjectGridSpringFactory"
      factory-method="getLocalPlatformTransactionManager"/>
</bean>

<bean id="Service" class="com.ibm.websphere.objectgrid.spring.test.TestService">
  <property name="txManager" ref="transactionManager"/>
</bean>
```

This shows the transactionManager bean being declared and wired in to the Service bean that will use Spring transactions. We will demonstrate this using annotations and this is the reason for the tx:annotation clause at the beginning.

- **Obtaining an ObjectGrid session:** A POJO that has methods managed by Spring can now obtain the ObjectGrid session for the current transaction using `Session s = txManager.getSession();`

This returns the session for the POJO to use. Beans participating in the same transaction will receive the same session when they call this method. Spring will automatically handle begin for the Session and also automatically invoke commit or rollback when necessary. You can obtain an ObjectGrid EntityManager also by simply calling `getEntityManager` from the Session object.

- **Setting the ObjectGrid instance for a thread:** A single Java Virtual Machine (JVM) can host many ObjectGrid instances. Each primary shard placed in a JVM has its own ObjectGrid instance. A JVM acting as a client to a remote ObjectGrid uses an ObjectGrid instance returned from the connect method's `ClientClusterContext` to interact with that Grid. Before invoking a method on a POJO using Spring transactions for ObjectGrid, the thread must be primed with the ObjectGrid instance to use. The `TransactionManager` instance has a method allowing a specific ObjectGrid instance to be specified. Once specified then any subsequent `txManager.getSession` calls will return Sessions for that ObjectGrid instance.

The following example shows a sample main for exercising this capability:

```
ClassPathXmlApplicationContext ctx = new ClassPathXmlApplicationContext(new String[]
    {"applicationContext.xml"});
SpringLocalTxManager txManager = (SpringLocalTxManager)ctx.getBean("transactionManager");
txManager.setObjectGridForThread(og);

ITestService s = (ITestService)ctx.getBean("Service");
s.initialize();
assertEquals(s.query(), "Billy");
s.update("Bobby");
assertEquals(s.query(), "Bobby");
System.out.println("Requires new test");
s.testRequiresNew(s);
assertEquals(s.query(), "1");
```

Here we use a Spring `ApplicationContext`. The `ApplicationContext` is used to obtain a reference to the `txManager` and specify an ObjectGrid to use on this thread. The code then obtains a reference to the service and invokes methods on it. Each method call at this level causes Spring to create a Session and do begin/commit calls around the method call. Any exceptions will cause a rollback.

- **SpringLocalTxManager interface:** The `SpringLocalTxManager` interface is implemented by the ObjectGrid Platform Transaction Manager and has all public interfaces. The methods on this interface are for selecting the ObjectGrid instance to use on a thread and obtaining a Session for the thread. Any POJOs using ObjectGrid local transactions should be injected with a reference to this manager instance and only a single instance need be created, that is, its scope should be singleton. This instance is created using a static method on `ObjectGridSpringFactory`. `getLocalPlatformTransactionManager()`.

Restriction: WebSphere eXtreme Scale does not support JTA or two-phase commit for various reasons mainly to do with scalability. Thus, except at a last single-phase participant, ObjectGrid does not interact in XA or JTA type global transactions. This platform manager is intended to make using local ObjectGrid transactions as easy as possible for Spring developers.

Spring managed extension beans

You can declare plain old Java objects (POJOs) to use as extension points in the `objectgrid.xml` file. If you name the beans and then specify the class name, eXtreme Scale normally creates instances of the specified class and uses those instances as the plug-in. WebSphere eXtreme Scale can now delegate to Spring to act as the bean factory for obtaining instances of these plug-in objects.

If an application uses Spring, POJOs have a requirement to be accessible to the rest of the application.

An application can register a Spring Bean Factory instance to use for an ObjectGrid specified by name. The application creates an instance of `BeanFactory` or a Spring application context and then registers it with ObjectGrid using the following static method:

```
void registerSpringBeanAdapterFactory(String objectGridName, Object springBeanFactory)
```

The previous method applies to the case when eXtreme Scale finds an extension bean whose `className` begins with the prefix `{spring}`. Such an extension bean, which could be an `ObjectTransformer`, `Loader`, `TransactionCallback`, and so on, uses the remainder of the name as a Spring Bean name. Then it obtains the bean instance using the Spring Bean Factory.

The eXtreme Scale deployment environment can also create a Spring bean factory from a default Spring XML configuration file. If no bean factory was registered for a given ObjectGrid, then your deployment searches for an XML file called `"/<ObjectGridName>_spring.xml"` automatically. For example, if your data grid is called `GRID`, then the XML file is called `"/GRID_spring.xml"` and appears in the class path in the root package. ObjectGrid constructs an `ApplicationContext` using the `"/<ObjectGridName>_spring.xml"` file and constructs beans from that bean factory.

The following is an example class name:

```
"{spring}MyLoaderBean"
```

Using the previous class name allows eXtreme Scale to use Spring to search for a bean named `"MyLoaderBean"`. You can specify Spring-managed POJOs for any extension point if the bean factory has been registered. The Spring extensions are in the `ogspring.jar` file. This JAR file must be on the class path for Spring support. If a J2EE application runs in WebSphere Application Server Network Deployment augmented with WebSphere Extended Deployment, then you must place the application on the application should place the `spring.jar` file and its associated files in the EAR modules. The `ogspring.jar` must also be placed in the same location.

Spring extension beans and namespace support

WebSphere eXtreme Scale provides a feature to declare plain old Java objects (POJOs) to use as extension points in the `objectgrid.xml` file and a way to name the beans and then specify the class name. Normally, instances of the specified class are created, and those objects are used as the plug-ins. Now, eXtreme Scale can delegate to Spring to obtain instances of these plug-in objects. If an application uses Spring then typically such POJOs have a requirement to be wired in to the rest of the application.

In some scenarios, you must use Spring to configure a plug-in, as in the following example:

```

<objectGrid name="Grid">
  <bean id="TransactionCallback" className="com.ibm.websphere.objectgrid.jpa.JPATxCallback">
    <property name="persistenceUnitName" type="java.lang.String" value="employeePU" />
  </bean>
  ...
</objectGrid>

```

The built-in TransactionCallback implementation, the `com.ibm.websphere.objectgrid.jpa.JPATxCallback` class, is configured as the TransactionCallback class. This class is configured with the **persistenceUnitName** property as shown in the previous example. The JPATxCallback class also has the JPAPropertyFactory attribute, which is of type `java.lang.Object`. The ObjectGrid XML configuration cannot support this type of configuration.

The eXtreme Scale Spring integration solves this problem by delegating the bean creation to the Spring framework. The revised configuration follows:

```

<objectGrid name="Grid">
  <bean id="TransactionCallback" className="{spring}jpaTxCallback"/>
  ...
</objectGrid>

```

The spring file for the "Grid" object contains the following information:

```

<bean id="jpaTxCallback" class="com.ibm.websphere.objectgrid.jpa.JPATxCallback" scope="shard">
  <property name="persistenceUnitName" value="employeeEMPU"/>
  <property name="JPAPropertyFactory" ref="jpaPropFactory"/>
</bean>

<bean id="jpaPropFactory" class="com.ibm.ws.objectgrid.jpa.plugins.
JPAPropFactoryImpl" scope="shard">
</bean>

```

Here, the TransactionCallback is specified as `{spring}jpaTxCallback`, and the `jpaTxCallback` and `jpaPropFactory` bean are configured in the spring file as shown in the previous example. The Spring configuration makes it possible to configure a JPAPropertyFactory bean as a parameter of the JPATxCallback object.

Default Spring bean factory

When eXtreme Scale finds a plug-in or an extension bean (such as an ObjectTransformer, Loader, TransactionCallback, and so on) with a classname value that begins with the prefix `{spring}`, then eXtreme Scale uses the remainder of the name as a Spring Bean name and obtain the bean instance using the Spring Bean Factory.

By default, if no bean factory was registered for a given ObjectGrid, then it tries to find an `ObjectGridName_spring.xml` file. For example, if your data grid is called "Grid" then the XML file is called `/Grid_spring.xml`. This file should be in the class path or in a META-INF directory which is in the class path. If this file is found, then eXtreme Scale constructs an ApplicationContext using that file and constructs beans from that bean factory.

Custom Spring bean factory

WebSphere eXtreme Scale also provides an ObjectGridSpringFactory API to register a Spring Bean Factory instance to use for a specific named ObjectGrid. This API registers an instance of BeanFactory with eXtreme Scale using the following static method:

```

void registerSpringBeanFactoryAdapter(String objectGridName, Object
springBeanFactory)

```

Namespace support

Since version 2.0, Spring has a mechanism for schema-based extensions to the basic Spring XML format for defining and configuring beans. ObjectGrid uses this new feature to define and configure ObjectGrid beans. With Spring XML schema extension, some of the built-in implementations of eXtreme Scale plug-ins and some ObjectGrid beans are predefined in the "objectgrid" namespace. When writing the Spring configuration files, you do not have to specify the full class name of the built-in implementations. Instead, you can reference the predefined beans.

Also, with the attributes of the beans defined in the XML schema, you are less likely to provide a wrong attribute name. XML validation based on the XML schema can catch these kind of errors earlier in the development cycle.

These beans defined in the XML schema extensions are:

- transactionManager
- register
- server
- catalog
- catalogServerProperties
- container
- JPALoader
- JPATxCallback
- JPAEntityLoader
- LRUEvictor
- LFUEvictor
- HashIndex

These beans are defined in the objectgrid.xsd XML schema. This XSD file is shipped as com/ibm/ws/objectgrid/spring/namespace/objectgrid.xsd file in the ogspring.jar file . For detailed descriptions of the XSD file and the beans defined in the XSD file, see the information about the Spring descriptor file in the *Administration Guide*.

Use the JPATxCallback example from the previous section. In the previous section, the JPATxCallback bean is configured as the following:

```
<bean id="jpaTxCallback" class="com.ibm.websphere.objectgrid.jpa.JPATxCallback" scope="shard">
  <property name="persistenceUnitName" value="employeeEMPU"/>
  <property name="JPAPropertyFactory" ref="jpaPropFactory"/>
</bean>

<bean id="jpaPropFactory" class="com.ibm.ws.objectgrid.jpa.plugins.JPAPropFactoryImpl" scope="shard">
</bean>
```

Using this namespace feature, the spring XML configuration can be written as the following:

```
<objectgrid:JPATxCallback id="jpaTxCallback" persistenceUnitName="employeeEMPU"
  jpaPropertyFactory="jpaPropFactory" />

<bean id="jpaPropFactory" class="com.ibm.ws.objectgrid.jpa.plugins.JPAPropFactoryImpl"
  scope="shard">
</bean>
```

Notice here that instead of specifying the com.ibm.websphere.objectgrid.jpa.JPATxCallback class as in the previous example, we directly use the pre-defined objectgrid:JPATxCallback bean. As you can see, this configuration is less verbose and more friendly to error checking.

For a description of working with Spring beans, consult “Starting a container server with Spring.”

Starting a container server with Spring

You can start a container server using Spring managed extension beans and namespace support.

About this task

With several XML files configured for Spring, you can start basic eXtreme Scale container servers.

Procedure

1. ObjectGrid XML file:

First of all, define a very simple ObjectGrid XML file which contains one ObjectGrid "Grid" and one map "Test". The ObjectGrid has an ObjectGridEventListener plug-in called "partitionListener", and the map "Test" has an Evictor plugged in called "testLRUEvictor". Notice both the ObjectGridEventListener plug-in and Evictor plug-in are configured using Spring as their names contain "{spring}".

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectgrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">
  <objectGrids>
    <objectGrid name="Grid">
      <bean id="ObjectGridEventListener" className="{spring}partitionListener" />
      <backingMap name="Test" pluginCollectionRef="test" />
    </objectGrid>
  </objectGrids>

  <backingMapPluginCollections>
    <backingMapPluginCollection id="test">
      <bean id="Evictor" className="{spring}testLRUEvictor"/>
    </backingMapPluginCollection>
  </backingMapPluginCollections>
</objectGridConfig>
```

2. ObjectGrid deployment XML file:

Now, create a simple ObjectGrid deployment XML file as follows. It partitions the ObjectGrid into 5 partitions, and no replica is required.

```
<?xml version="1.0" encoding="UTF-8"?>
<deploymentPolicy xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/deploymentPolicy ../deploymentPolicy.xsd"
  xmlns="http://ibm.com/ws/objectgrid/deploymentPolicy">
  <objectgridDeployment objectgridName="Grid">
    <mapSet name="mapSet" numInitialContainers="1" numberOfPartitions="5" minSyncReplicas="0"
      maxSyncReplicas="1" maxAsyncReplicas="0">
      <map ref="Test"/>
    </mapSet>
  </objectgridDeployment>
</deploymentPolicy>
```

3. ObjectGrid Spring XML file:

Now we will use both ObjectGrid Spring managed extension beans and namespace support features to configure the ObjectGrid beans. The spring xml file is named Grid_spring.xml. Notice two schemas are included in the XML file: spring-beans-2.0.xsd is for using the Spring managed beans, and objectgrid.xsd is for using the beans predefined in the objectgrid namespace.

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:aop="http://www.springframework.org/schema/aop"
```

```

    xmlns:tx="http://www.springframework.org/schema/tx"
    xmlns:objectgrid="http://www.ibm.com/schema/objectgrid"
    xsi:schemaLocation="
    http://www.ibm.com/schema/objectgrid
    http://www.ibm.com/schema/objectgrid/objectgrid.xsd
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans-2.0.xsd">

    <objectgrid:register id="ogregister" gridname="Grid"/>

    <objectgrid:server id="server" isCatalog="true" name="server">
        <objectgrid:catalog host="localhost" port="2809"/>
    </objectgrid:server>

    <objectgrid:container id="container"
    objectgridxml="com/ibm/ws/objectgrid/test/springshard/objectgrid.xml"
    deploymentxml="com/ibm/ws/objectgrid/test/springshard/deployment.xml"
    server="server"/>

    <objectgrid:LRUEvictor id="testLRUEvictor" numberOfLRUQueues="31"/>

    <bean id="partitionListener"
    class="com.ibm.websphere.objectgrid.springshard.ShardListener" scope="shard"/>
</beans>

```

There were six beans defined in this spring XML file:

- a. *objectgrid:register*: This register the default bean factory for the ObjectGrid "Grid".
- b. *objectgrid:server*: This defines an ObjectGrid server with name "server". This server will also provide catalog service since it has an objectgrid:catalog bean nested in it.
- c. *objectgrid:catalog*: This defines an ObjectGrid catalog service endpoint, which is set to "localhost:2809".
- d. *objectgrid:container*: This defines an ObjectGrid container with specified objectgrid XML file and deployment XML file as we discussed before. The server property specifies which server this container is hosted in.
- e. *objectgrid:LRUEvictor*: This defines an LRUEvictor with the number of LRU queues to use set to 31.
- f. *bean partitionListener*: This defines a ShardListener plug-in. You must provide an implementation for this plug-in, so it cannot use the pre-defined beans. Also this scope of the bean is set to "shard", which means there is only one instance of this ShardListener per ObjectGrid shard.

4. Starting the server:

The snippet below starts the ObjectGrid server, which hosts both the container service and the catalog service. As we can see, the only method we need to call to start the server is to get a bean "container" from the bean factory. This simplifies the programming complexity by moving most of the logic into Spring configuration.

```

public class ShardServer extends TestCase
{
    Container container;
    org.springframework.beans.factory.BeanFactory bf;

    public void startServer(String cep)
    {
        try
        {
            bf = new org.springframework.context.support.ClassPathXmlApplicationContext(
                "/com/ibm/ws/objectgrid/test/springshard/Grid_spring.xml", ShardServer.class);
            container = (Container)bf.getBean("container");
        }
        catch(Exception e)
        {

```

```
        throw new ObjectGridRuntimeException("Cannot start OG container", e);
    }
}

public void stopServer()
{
    if(container != null)
        container.teardown();
}
}
```


Chapter 8. Programming for security

Use programming interfaces to handle various aspects of security in a WebSphere eXtreme Scale environment.

Security API

WebSphere eXtreme Scale adopts an open security architecture. It provides a basic security framework for authentication, authorization, and transport security, and requires users to implement plug-ins to complete the security infrastructure.

The following image shows the basic flow of client authentication and authorization for an eXtreme Scale server.

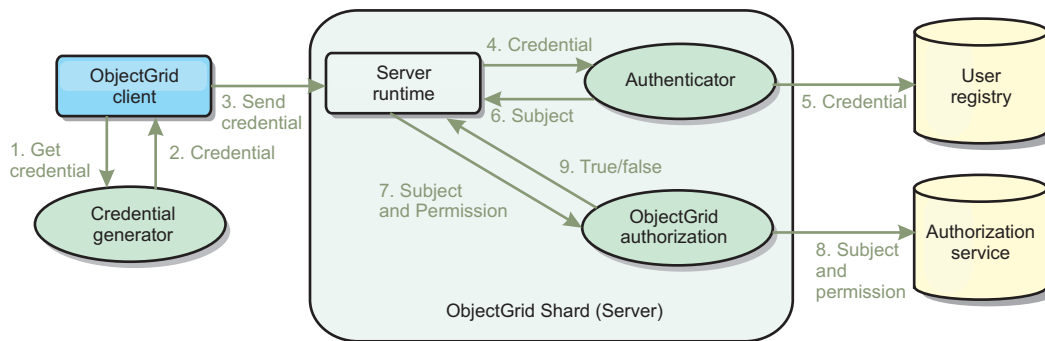


Figure 17. Flow of client authentication and authorization

The authentication flow and authorization flow are as follows.

Authentication flow

1. The authentication flow starts with an eXtreme Scale client getting a credential. This is done by the `com.ibm.websphere.objectgrid.security.plugins.CredentialGenerator` plug-in.
2. A `CredentialGenerator` object knows how to generate a valid client credential, for example, a user ID and password pair, Kerberos ticket, and so on. This generated credential is sent back to the client.
3. After the client retrieves the `Credential` object using the `CredentialGenerator` object, this `Credential` object is sent along with the eXtreme Scale request to the eXtreme Scale server.
4. The eXtreme Scale server authenticates the `Credential` object before processing the eXtreme Scale request. Then the server uses the `Authenticator` plug-in to authenticate the `Credential` object.
5. The `Authenticator` plug-in represents an interface to the user registry, for example, a Lightweight Directory Access Protocol (LDAP) server or an operating system user registry. The `Authenticator` consults the user registry and makes authentication decisions.
6. If the authentication is successful, a `Subject` object is returned to represent this client.

Authorization flow

WebSphere eXtreme Scale adopts a permission-based authorization mechanism, and has different permission categories represented by different permission classes. For example, a `com.ibm.websphere.objectgrid.security.MapPermission` object represents permissions to read, write, insert, invalidate, and remove the data entries in an `ObjectMap`. Because WebSphere eXtreme Scale supports Java Authentication and Authorization Service (JAAS) authorization out-of-box, you can use JAAS to handle authorization by providing authorization policies.

Also, eXtreme Scale supports custom authorizations. Custom authorizations are plugged in by the plug-in

`com.ibm.websphere.objectgrid.security.plugins.ObjectGridAuthorization`. The flow of the customer authorization is as follows.

7. The server runtime sends the Subject object and the required permission to the authorization plug-in.
8. The authorization plug-in consults the Authorization service and makes an authorization decision. If permission is granted for this Subject object, a value of `true` is returned, otherwise `false` is returned.
9. This authorization decision, `true` or `false`, is returned to the server runtime.

Security implementation

The topics in this section discuss how to program a secure WebSphere eXtreme Scale deployment and how to program the plug-in implementations. The section is organized based on the various security features. In each subtopic, you will learn about relevant plug-ins and how to implement the plug-ins. In the authentication section, you will see how to connect to a secure WebSphere eXtreme Scale deployment environment.

Client Authentication: The client authentication topic describes how a WebSphere eXtreme Scale client gets a credential and how a server authenticates the client. It will also discuss how a WebSphere eXtreme Scale client connects to a secure WebSphere eXtreme Scale server.

Authorization: The authorization topic explains how to use the `ObjectGridAuthorization` to do customer authorization besides JAAS authorization.

Grid Authentication: The data grid authentication topic discusses how you can use `SecureTokenManager` to securely transport server secrets.

Java Management Extensions (JMX) programming: When the WebSphere eXtreme Scale server is secured, the JMX client might need to send a JMX credential to the server.

Client authentication programming

For authentication, WebSphere eXtreme Scale provides a runtime to send the credential from the client to the server side, and then calls the authenticator plug-in to authenticate the users.

WebSphere eXtreme Scale requires you to implement the following plug-ins to complete the authentication.

- **Credential:** A `Credential` represents a client credential, such as a user ID and password pair.
- **CredentialGenerator:** A `CredentialGenerator` represents a credential factory to generate the credential.

- Authenticator: An Authenticator authenticates the client credential and retrieves client information.

Credential and CredentialGenerator plug-ins

When an eXtreme Scale client connects to a server that requires authentication, the client is required to provide a client credential. A client credential is represented by a `com.ibm.websphere.objectgrid.security.plugins.Credential` interface. A client credential can be a user name and password pair, a Kerberos ticket, a client certificate, or data in any format that the client and server agree upon. See the information about the Credential API in the API documentation for more details. This interface explicitly defines the `equals(Object)` and `hashCode` methods. These two methods are important because the authenticated Subject objects are cached by using the Credential object as the key on the server side. WebSphere eXtreme Scale also provides a plug-in to generate a credential. This plug-in is represented by the `com.ibm.websphere.objectgrid.security.plugins.CredentialGenerator` interface and is useful when the credential is can expire. In this case, the `getCredential` method is called to renew a credential.

The Credential interface explicitly defines the `equals(Object)` and `hashCode` methods. These two methods are important because the authenticated Subject objects are cached by using the Credential object as the key on the server side.

You may also use the provided plug-in to generate a credential. This plug-in is represented by the `com.ibm.websphere.objectgrid.security.plugins.CredentialGenerator` interface, and is useful when the credential can expire. In this case, the `getCredential` method is called to renew a credential. See the API documentation for more details.

There are three provided default implementations for the Credential interfaces:

- The `com.ibm.websphere.objectgrid.security.plugins.builtins.UserPasswordCredential` implementation, which contains a user ID and password pair.
- The `com.ibm.websphere.objectgrid.security.plugins.builtins.WSTokenCredential` implementation, which contains WebSphere Application Server-specific authentication and authorization tokens. These tokens can be used to propagate the security attributes across the application servers in the same security domain.

WebSphere eXtreme Scale also provides a plug-in to generate a credential. This plug-in is represented by the `com.ibm.websphere.objectgrid.security.plugins.CredentialGenerator` interface. WebSphere eXtreme Scale provides two default built-in implementations:

- The `com.ibm.websphere.objectgrid.security.plugins.builtins.UserPasswordCredentialGenerator` constructor takes a user ID and a password. When the `getCredential` method is called, it returns a `UserPasswordCredential` object that contains the user ID and password.
- The `com.ibm.websphere.objectgrid.security.plugins.builtins.WSTokenCredentialGenerator` represents a credential (security token) generator when running in WebSphere Application Server. When the `getCredential` method is called, the Subject that is associated with the current thread is retrieved. Then the security information in this Subject object is converted into a `WSTokenCredential` object. You can specify whether to retrieve a `runAs` subject or a caller subject from the thread by using the constant

WSTokenCredentialGenerator.RUN_AS_SUBJECT or
WSTokenCredentialGenerator.CALLER_SUBJECT.

UserPasswordCredential and UserPasswordCredentialGenerator

For testing purposes, WebSphere eXtreme Scale provides the following plug-in implementations:

1. `com.ibm.websphere.objectgrid.security.plugins.builtins.UserPasswordCredential`
2. `com.ibm.websphere.objectgrid.security.plugins.builtins.UserPasswordCredentialGenerator`

The user password credential stores a user ID and password. The user password credential generator then contains this user ID and password.

The following example code shows how to implement these two plug-ins.

```
UserPasswordCredential.java
// This sample program is provided AS IS and may be used, executed, copied and modified
// without royalty payment by customer
// (a) for its own instruction and study,
// (b) in order to develop applications designed to run with an IBM WebSphere product,
// either for customer's own internal use or for redistribution by customer, as part of such an
// application, in customer's own products.
// Licensed Materials - Property of IBM
// 5724-J34 © COPYRIGHT International Business Machines Corp. 2007
package com.ibm.websphere.objectgrid.security.plugins.builtins;

import com.ibm.websphere.objectgrid.security.plugins.Credential;

/**
 * This class represents a credential containing a user ID and password.
 *
 * @ibm-api
 * @since WAS XD 6.0.1
 *
 * @see Credential
 * @see UserPasswordCredentialGenerator#getCredential()
 */
public class UserPasswordCredential implements Credential {

    private static final long serialVersionUID = 1409044825541007228L;

    private String ivUserName;

    private String ivPassword;

    /**
     * Creates a UserPasswordCredential with the specified user name and
     * password.
     *
     * @param userName the user name for this credential
     * @param password the password for this credential
     *
     * @throws IllegalArgumentException if userName or password is <code>null</code>
     */
    public UserPasswordCredential(String userName, String password) {
        super();
        if (userName == null || password == null) {
            throw new IllegalArgumentException("User name and password cannot be null.");
        }
        this.ivUserName = userName;
        this.ivPassword = password;
    }

    /**
     * Gets the user name for this credential.
     *
     * @return the user name argument that was passed to the constructor
     * or the <code>setUserName(String)</code>
     * method of this class
     *
     * @see #setUserName(String)
     */
    public String getUserName() {
        return ivUserName;
    }

    /**
     * Sets the user name for this credential.
     *
     * @param userName the user name to set.
     */
}
```



```

    *
    * @throws IllegalArgumentException if userName is <code>null</code>
    */
    public void setUsername(String userName) {
        if (userName == null) {
            throw new IllegalArgumentException("User name cannot be null.");
        }
        this.ivUserName = userName;
    }

    /**
     * Gets the password for this credential.
     *
     * @return the password argument that was passed to the constructor
     *         or the <code>setPassword(String)</code>
     *         method of this class
     *
     * @see #setPassword(String)
     */
    public String getPassword() {
        return ivPassword;
    }

    /**
     * Sets the password for this credential.
     *
     * @param password the password to set.
     *
     * @throws IllegalArgumentException if password is <code>null</code>
     */
    public void setPassword(String password) {
        if (password == null) {
            throw new IllegalArgumentException("Password cannot be null.");
        }
        this.ivPassword = password;
    }

    /**
     * Checks two UserPasswordCredential objects for equality.
     * <p>
     * Two UserPasswordCredential objects are equal if and only if their user names
     * and passwords are equal.
     *
     * @param o the object we are testing for equality with this object.
     *
     * @return <code>true</code> if both UserPasswordCredential objects are equivalent.
     *
     * @see Credential#equals(Object)
     */
    public boolean equals(Object o) {
        if (this == o) {
            return true;
        }
        if (o instanceof UserPasswordCredential) {
            UserPasswordCredential other = (UserPasswordCredential) o;
            return other.ivPassword.equals(ivPassword) && other.ivUserName.equals(ivUserName);
        }
        return false;
    }

    /**
     * Returns the hashCode of the UserPasswordCredential object.
     *
     * @return the hash code of this object
     *
     * @see Credential#hashCode()
     */
    public int hashCode() {
        return ivUserName.hashCode() + ivPassword.hashCode();
    }
}

```

UserPasswordCredentialGenerator.java

```

// This sample program is provided AS IS and may be used, executed, copied and modified
// without royalty payment by customer
// (a) for its own instruction and study,
// (b) in order to develop applications designed to run with an IBM WebSphere product,
// either for customer's own internal use or for redistribution by customer, as part of such an
// application, in customer's own products.
// Licensed Materials - Property of IBM
// 5724-J34 © COPYRIGHT International Business Machines Corp. 2007
package com.ibm.websphere.objectgrid.security.plugins.builtins;

```

```
import java.util.StringTokenizer;
```

```
import com.ibm.websphere.objectgrid.security.plugins.Credential;
import com.ibm.websphere.objectgrid.security.plugins.CredentialGenerator;
```

```

/**
 * This credential generator creates <code>UserPasswordCredential</code> objects.
 * <p>

```

```

* UserPasswordCredentialGenerator has a one to one relationship with
* UserPasswordCredential because it can only create a UserPasswordCredential
* representing one identity.
*
* @since WAS XD 6.0.1
* @ibm-api
*
* @see CredentialGenerator
* @see UserPasswordCredential
*/
public class UserPasswordCredentialGenerator implements CredentialGenerator {

    private String ivUser;

    private String ivPwd;

    /**
     * Creates a UserPasswordCredentialGenerator with no user name or password.
     *
     * @see #setProperties(String)
     */
    public UserPasswordCredentialGenerator() {
        super();
    }

    /**
     * Creates a UserPasswordCredentialGenerator with a specified user name and
     * password
     *
     * @param user the user name
     * @param pwd the password
     */
    public UserPasswordCredentialGenerator(String user, String pwd) {
        ivUser = user;
        ivPwd = pwd;
    }

    /**
     * Creates a new <code>UserPasswordCredential</code> object using this
     * object's user name and password.
     *
     * @return a new <code>UserPasswordCredential</code> instance
     *
     * @see CredentialGenerator#getCredential()
     * @see UserPasswordCredential
     */
    public Credential getCredential() {
        return new UserPasswordCredential(ivUser, ivPwd);
    }

    /**
     * Gets the password for this credential generator.
     *
     * @return the password argument that was passed to the constructor
     */
    public String getPassword() {
        return ivPwd;
    }

    /**
     * Gets the user name for this credential.
     *
     * @return the user argument that was passed to the constructor
     *         of this class
     */
    public String getUsername() {
        return ivUser;
    }

    /**
     * Sets additional properties namely a user name and password.
     *
     * @param properties a properties string with a user name and
     *                  a password separated by a blank.
     *
     * @throws IllegalArgumentException if the format is not valid
     */
    public void setProperties(String properties) {
        StringTokenizer token = new StringTokenizer(properties, " ");
        if (token.countTokens() != 2) {
            throw new IllegalArgumentException(
                "The properties should have a user name and password and separated by a blank.");
        }

        ivUser = token.nextToken();
        ivPwd = token.nextToken();
    }

    /**
     * Checks two UserPasswordCredentialGenerator objects for equality.
     *
     * <p>
     * Two UserPasswordCredentialGenerator objects are equal if and only if
     * their user names and passwords are equal.
     */

```

```

*
* @param obj the object we are testing for equality with this object.
*
* @return <code>true</code> if both UserPasswordCredentialGenerator objects
*         are equivalent.
*/
public boolean equals(Object obj) {
    if (obj == this) {
        return true;
    }

    if (obj != null && obj instanceof UserPasswordCredentialGenerator) {
        UserPasswordCredentialGenerator other = (UserPasswordCredentialGenerator) obj;

        boolean bothUserNull = false;
        boolean bothPwdNull = false;

        if (ivUser == null) {
            if (other.ivUser == null) {
                bothUserNull = true;
            } else {
                return false;
            }
        }

        if (ivPwd == null) {
            if (other.ivPwd == null) {
                bothPwdNull = true;
            } else {
                return false;
            }
        }

        return (bothUserNull || ivUser.equals(other.ivUser)) && (bothPwdNull || ivPwd.equals(other.ivPwd));
    }

    return false;
}

/**
 * Returns the hashCode of the UserPasswordCredentialGenerator object.
 *
 * @return the hash code of this object
 */
public int hashCode() {
    return ivUser.hashCode() + ivPwd.hashCode();
}
}

```

The `UserPasswordCredential` class contains two attributes: user name and password. The `UserPasswordCredentialGenerator` serves as a factory that contains the `UserPasswordCredential` objects.

WSTokenCredential and WSTokenCredentialGenerator

When the WebSphere eXtreme Scale clients and servers are all deployed in WebSphere Application Server, the client application can use these two built-in implementations when the following conditions are satisfied:

1. WebSphere Application Server global security is turned on.
2. All WebSphere eXtreme Scale clients and servers are running in WebSphere Application Server Java virtual machines.
3. The application servers are in the same security domain.
4. The client is already authenticated in WebSphere Application Server.

In this situation, the client can use the `com.ibm.websphere.objectgrid.security.plugins.builtins.WSTokenCredentialGenerator` class to generate a credential. The server uses the `WSAuthenticator` implementation class to authenticate the credential.

This scenario takes advantage of the fact that the eXtreme Scale client has already been authenticated. Because the application servers that have the servers are in the same security domain as the application servers that house the clients, the security

tokens can be propagated from the client to the server so that the same user registry does not need to be authenticated again.

Note: Do not assume that a `CredentialGenerator` always generates the same credential. For an expirable and refreshable credential, the `CredentialGenerator` should be able to generate the latest valid credential to make sure the authentication succeeds. One example is using the Kerberos ticket as a `Credential` object. When the Kerberos ticket refreshes, the `CredentialGenerator` should retrieve the refreshed ticket when `CredentialGenerator.getCredential` is called.

Authenticator plug-in

After the eXtreme Scale client retrieves the `Credential` object using the `CredentialGenerator` object, this client `Credential` object is sent along with the client request to the eXtreme Scale server. The server authenticates the `Credential` object before processing the request. If the `Credential` object is authenticated successfully, a `Subject` object is returned to represent this client.

This `Subject` object is then cached, and it expires after its lifetime reaches the session timeout value. The login session timeout value can be set by using the `loginSessionExpirationTime` property in the cluster XML file. For example, setting `loginSessionExpirationTime="300"` makes the `Subject` object expire in 300 seconds.

This `Subject` object is then used for authorizing the request, which is shown later. An eXtreme Scale server uses the Authenticator plug-in to authenticate the `Credential` object. See the information about the Authenticator in the API documentation for more details.

The Authenticator plug-in is where the eXtreme Scale runtime authenticates the `Credential` object from the client user registry, for example, a Lightweight Directory Access Protocol (LDAP) server.

WebSphere eXtreme Scale does not provide an immediately available user registry configuration. The configuration and management of user registry is left outside of WebSphere eXtreme Scale for simplicity and flexibility. This plug-in implements connecting and authenticating to the user registry. For example, an Authenticator implementation extracts the user ID and password from the credential, uses them to connect and validate to an LDAP server, and creates a `Subject` object as a result of the authentication. The implementation might use JAAS login modules. A `Subject` object is returned as a result of authentication.

Notice that this method creates two exceptions: `InvalidCredentialException` and `ExpiredCredentialException`. The `InvalidCredentialException` exception indicates that the credential is not valid. The `ExpiredCredentialException` exception indicates that the credential expired. If one of these two exceptions result from the `authenticate` method, the exceptions are sent back to the client. However, the client runtime handles these two exceptions differently:

- If the error is an `InvalidCredentialException` exception, the client run time displays this exception. Your application must handle the exception. You can correct the `CredentialGenerator`, for example, and then retry the operation.
- If the error is an `ExpiredCredentialException` exception, and the retry count is not 0, the client run time calls the `CredentialGenerator.getCredential` method again, and sends the new `Credential` object to the server. If the new credential authentication succeeds, the server processes the request. If the new credential authentication fails, the exception is sent back to the client. If the number of

authentication retries reaches the supported value and the client still gets an `ExpiredCredentialException` exception, the `ExpiredCredentialException` exception results. Your application must handle the error.

The `Authenticator` interface provides great flexibility. You can implement the `Authenticator` interface in your own specific way. For example, you can implement this interface to support two different user registries.

WebSphere eXtreme Scale provides sample authenticator plug-in implementations. Except for the WebSphere Application Server authenticator plug-in, the other implementations are only samples for testing purposes.

KeyStoreLoginAuthenticator

This example uses an eXtreme Scale built-in implementation: `KeyStoreLoginAuthenticator`, which is for testing and sample purposes (a key store is a simple user registry and should not be used for a production environment). Again, the class is displayed to further demonstrate how to implement an authenticator.

```
KeyStoreLoginAuthenticator.java
// This sample program is provided AS IS and may be used, executed, copied and modified
// without royalty payment by customer
// (a) for its own instruction and study,
// (b) in order to develop applications designed to run with an IBM WebSphere product,
// either for customer's own internal use or for redistribution by customer, as part of such an
// application, in customer's own products.
// Licensed Materials - Property of IBM
// 5724-J34 © COPYRIGHT International Business Machines Corp. 2007

package com.ibm.websphere.objectgrid.security.plugins.builtins;

import javax.security.auth.Subject;
import javax.security.auth.login.LoginContext;
import javax.security.auth.login.LoginException;

import com.ibm.websphere.objectgrid.security.plugins.Authenticator;
import com.ibm.websphere.objectgrid.security.plugins.Credential;
import com.ibm.websphere.objectgrid.security.plugins.ExpiredCredentialException;
import com.ibm.websphere.objectgrid.security.plugins.InvalidCredentialException;
import com.ibm.ws.objectgrid.Constants;
import com.ibm.ws.objectgrid.ObjectGridManagerImpl;
import com.ibm.ws.objectgrid.security.auth.callback.UserPasswordCallbackHandlerImpl;

/**
 * This class is an implementation of the <code>Authenticator</code> interface
 * when a user name and password are used as a credential.
 * <p>
 * When user ID and password authentication is used, the credential passed to the
 * <code>authenticate(Credential)</code> method is a UserPasswordCredential object.
 * <p>
 * This implementation will use a <code>KeyStoreLoginModule</code> to authenticate
 * the user into the key store using the JAAS login module "KeyStoreLogin". The key
 * store can be configured as an option to the <code>KeyStoreLoginModule</code>
 * class. Please see the <code>KeyStoreLoginModule</code> class for more details
 * about how to set up the JAAS login configuration file.
 * <p>
 * This class is only for sample and quick testing purpose. Users should
 * write your own Authenticator implementation which can fit better into
 * the environment.
 *
 * @ibm-api
 * @since WAS XD 6.0.1
 *
 * @see Authenticator
 * @see KeyStoreLoginModule
 * @see UserPasswordCredential
 */
public class KeyStoreLoginAuthenticator implements Authenticator {

    /**
     * Creates a new KeyStoreLoginAuthenticator.
     */
    public KeyStoreLoginAuthenticator() {
        super();
    }

    /**
     * Authenticates a <code>UserPasswordCredential</code>.
     * <p>

```

```

    * Uses the user name and password from the specified UserPasswordCredential
    * to login to the KeyStoreLoginModule named "KeyStoreLogin".
    *
    * @throws InvalidCredentialException if credential isn't a
    *       UserPasswordCredential or some error occurs during processing
    *       of the supplied UserPasswordCredential
    *
    * @throws ExpiredCredentialException if credential is expired. This exception
    *       is not used by this implementation
    *
    * @see Authenticator#authenticate(Credential)
    * @see KeyStoreLoginModule
    */
    public Subject authenticate(Credential credential) throws InvalidCredentialException,
    ExpiredCredentialException {

        if (credential == null) {
            throw new InvalidCredentialException("Supplied credential is null");
        }

        if (! (credential instanceof UserPasswordCredential) ) {
            throw new InvalidCredentialException("Supplied credential is not a UserPasswordCredential");
        }

        UserPasswordCredential cred = (UserPasswordCredential) credential;
        LoginContext lc = null;
        try {
            lc = new LoginContext("KeyStoreLogin",
                new UserPasswordCallbackHandlerImpl(cred.getUserName(), cred.getPassword().toCharArray()));

            lc.login();

            Subject subject = lc.getSubject();

            return subject;
        }
        catch (LoginException le) {
            throw new InvalidCredentialException(le);
        }
        catch (IllegalArgumentException ile) {
            throw new InvalidCredentialException(ile);
        }
    }
}

```

KeyStoreLoginModule.java

```

// This sample program is provided AS IS and may be used, executed, copied and modified
// without royalty payment by customer
// (a) for its own instruction and study,
// (b) in order to develop applications designed to run with an IBM WebSphere product,
// either for customer's own internal use or for redistribution by customer, as part of such an
// application, in customer's own products.
// Licensed Materials - Property of IBM
// 5724-J34 © COPYRIGHT International Business Machines Corp. 2007
package com.ibm.websphere.objectgrid.security.plugins.builtins;

import java.io.File;
import java.io.FileInputStream;
import java.security.KeyStore;
import java.security.KeyStoreException;
import java.security.NoSuchAlgorithmException;
import java.security.PrivateKey;
import java.security.UnrecoverableKeyException;
import java.security.cert.Certificate;
import java.security.cert.CertificateException;
import java.security.cert.CertificateFactory;
import java.security.cert.X509Certificate;
import java.util.Arrays;
import java.util.HashSet;
import java.util.Map;
import java.util.Set;

import javax.security.auth.Subject;
import javax.security.auth.callback.Callback;
import javax.security.auth.callback.CallbackHandler;
import javax.security.auth.callback.NameCallback;
import javax.security.auth.callback.PasswordCallback;
import javax.security.auth.login.LoginException;
import javax.security.auth.spi.LoginModule;
import javax.security.auth.x500.X500Principal;
import javax.security.auth.x500.X500PrivateCredential;

import com.ibm.websphere.objectgrid.ObjectGridRuntimeException;
import com.ibm.ws.objectgrid.Constants;
import com.ibm.ws.objectgrid.ObjectGridManagerImpl;
import com.ibm.ws.objectgrid.util.ObjectGridUtil;

/**
 * A KeyStoreLoginModule is keystore authentication login module based on
 * JAAS authentication.
 * <p>
 * A login configuration should provide an option "<code>keyStoreFile</code>" to

```

```

* indicate where the keystore file is located. If the <code>keyStoreFile</code>
* value contains a system property in the form, <code>${system.property}</code>,
* it will be expanded to the value of the system property.
* <p>
* If an option "<code>keyStoreFile</code>" is not provided, the default keystore
* file name is <code>${java.home}/.keystore</code>.
* <p>
* Here is a Login module configuration example:
* <pre><code>
*   KeyStoreLogin {
*       com.ibm.websphere.objectgrid.security.plugins.builtins.KeystoreLoginModule required
*       keyStoreFile="${user.dir}/${}security/${}.keystore";
*   };
* </code></pre>
*
* @ibm-api
* @since WAS XD 6.0.1
*
* @see LoginModule
*/
public class KeyStoreLoginModule implements LoginModule {

    private static final String CLASS_NAME = KeyStoreLoginModule.class.getName();

    /**
     * Key store file property name
     */
    public static final String KEY_STORE_FILE_PROPERTY_NAME = "keyStoreFile";

    /**
     * Key store type. Only JKS is supported
     */
    public static final String KEYSTORE_TYPE = "JKS";

    /**
     * The default key store file name
     */
    public static final String DEFAULT_KEY_STORE_FILE = "${java.home}/${}.keystore";

    private CallbackHandler handler;

    private Subject subject;

    private boolean debug = false;

    private Set principals = new HashSet();

    private Set publicCreds = new HashSet();

    private Set privateCreds = new HashSet();

    protected KeyStore keyStore;

    /**
     * Creates a new KeyStoreLoginModule.
     */
    public KeyStoreLoginModule() {
    }

    /**
     * Initializes the login module.
     *
     * @see LoginModule#initialize(Subject, CallbackHandler, Map, Map)
     */
    public void initialize(Subject sub, CallbackHandler callbackHandler,
        Map mapSharedState, Map mapOptions) {

        // initialize any configured options
        debug = "true".equalsIgnoreCase((String) mapOptions.get("debug"));

        if (sub == null)
            throw new IllegalArgumentException("Subject is not specified");

        if (callbackHandler == null)
            throw new IllegalArgumentException(
                "CallbackHandler is not specified");

        // Get the key store path
        String sKeyStorePath = (String) mapOptions
            .get(KEY_STORE_FILE_PROPERTY_NAME);

        // If there is no key store path, the default one is the .keystore
        // file in the java home directory
        if (sKeyStorePath == null) {
            sKeyStorePath = DEFAULT_KEY_STORE_FILE;
        }

        // Replace the system environment variable
        sKeyStorePath = ObjectGridUtil.replaceVar(sKeyStorePath);

        File fileKeyStore = new File(sKeyStorePath);

```

```

try {
    KeyStore store = KeyStore.getInstance("JKS");
    store.load(new FileInputStream(fileKeyStore), null);

    // Save the key store
    keyStore = store;

    if (debug) {
        System.out.println("[KeyStoreLoginModule] initialize: Successfully loaded key store");
    }
}
catch (Exception e) {
    ObjectGridRuntimeException re = new ObjectGridRuntimeException(
        "Failed to load keystore: " + fileKeyStore.getAbsolutePath());
    re.initCause(e);
    if (debug) {
        System.out.println("[KeyStoreLoginModule] initialize: Key store loading failed with exception "
            + e.getMessage());
    }
}

this.subject = sub;
this.handler = callbackHandler;
}

/**
 * Authenticates a user based on the keystore file.
 *
 * @see LoginModule#login()
 */
public boolean login() throws LoginException {

    if (debug) {
        System.out.println("[KeyStoreLoginModule] login: entry");
    }

    String name = null;
    char pwd[] = null;

    if (keyStore == null || subject == null || handler == null) {
        throw new LoginException("Module initialization failed");
    }

    NameCallback nameCallback = new NameCallback("Username:");
    PasswordCallback pwdCallback = new PasswordCallback("Password:", false);

    try {
        handler.handle(new Callback[] { nameCallback, pwdCallback });
    }
    catch (Exception e) {
        throw new LoginException("Callback failed: " + e);
    }

    name = nameCallback.getName();
    char[] tempPwd = pwdCallback.getPassword();

    if (tempPwd == null) {
        // treat a NULL password as an empty password
        tempPwd = new char[0];
    }
    pwd = new char[tempPwd.length];
    System.arraycopy(tempPwd, 0, pwd, 0, tempPwd.length);

    pwdCallback.clearPassword();

    if (debug) {
        System.out.println("[KeyStoreLoginModule] login: "
            + "user entered user name: " + name);
    }

    // Validate the user name and password
    try {
        validate(name, pwd);
    }
    catch (SecurityException se) {
        principals.clear();
        publicCreds.clear();
        privateCreds.clear();
        LoginException le = new LoginException(
            "Exception encountered during login");
        le.initCause(se);

        throw le;
    }

    if (debug) {
        System.out.println("[KeyStoreLoginModule] login: exit");
    }
    return true;
}

```



```

/**
 * Indicates the user is accepted.
 * <p>
 * This method is called only if the user is authenticated by all modules in
 * the login configuration file. The principal objects will be added to the
 * stored subject.
 *
 * @return false if for some reason the principals cannot be added; true
 *         otherwise
 *
 * @exception LoginException
 *         LoginException is thrown if the subject is readonly or if
 *         any unrecoverable exceptions is encountered.
 *
 * @see LoginModule#commit()
 */
public boolean commit() throws LoginException {
    if (debug) {
        System.out.println("[KeyStoreLoginModule] commit: entry");
    }

    if (principals.isEmpty()) {
        throw new IllegalStateException("Commit is called out of sequence");
    }

    if (subject.isReadOnly()) {
        throw new LoginException("Subject is ReadOnly");
    }

    subject.getPrincipals().addAll(principals);
    subject.getPublicCredentials().addAll(publicCreds);
    subject.getPrivateCredentials().addAll(privateCreds);

    principals.clear();
    publicCreds.clear();
    privateCreds.clear();

    if (debug) {
        System.out.println("[KeyStoreLoginModule] commit: exit");
    }
    return true;
}

/**
 * Indicates the user is not accepted
 *
 * @see LoginModule#abort()
 */
public boolean abort() throws LoginException {
    boolean b = logout();
    return b;
}

/**
 * Logs the user out. Clear all the maps.
 *
 * @see LoginModule#logout()
 */
public boolean logout() throws LoginException {

    // Clear the instance variables
    principals.clear();
    publicCreds.clear();
    privateCreds.clear();

    // clear maps in the subject
    if (!subject.isReadOnly()) {
        if (subject.getPrincipals() != null) {
            subject.getPrincipals().clear();
        }

        if (subject.getPublicCredentials() != null) {
            subject.getPublicCredentials().clear();
        }

        if (subject.getPrivateCredentials() != null) {
            subject.getPrivateCredentials().clear();
        }
    }
    return true;
}

/**
 * Validates the user name and password based on the keystore.
 *
 * @param userName user name
 * @param password password
 * @throws SecurityException if any exceptions encountered
 */

```

```

private void validate(String userName, char password[])
    throws SecurityException {

    PrivateKey privateKey = null;

    // Get the private key from the keystore
    try {
        privateKey = (PrivateKey) keyStore.getKey(userName, password);
    }
    catch (NoSuchAlgorithmException nsae) {
        SecurityException se = new SecurityException();
        se.initCause(nsae);
        throw se;
    }
    catch (KeyStoreException kse) {
        SecurityException se = new SecurityException();
        se.initCause(kse);
        throw se;
    }
    catch (UnrecoverableKeyException uke) {
        SecurityException se = new SecurityException();
        se.initCause(uke);
        throw se;
    }

    if (privateKey == null) {
        throw new SecurityException("Invalid name: " + userName);
    }

    // Check the certificates
    Certificate certs[] = null;
    try {
        certs = keyStore.getCertificateChain(userName);
    }
    catch (KeyStoreException kse) {
        SecurityException se = new SecurityException();
        se.initCause(kse);
        throw se;
    }

    if (debug) {
        System.out.println(" Print out the certificates:");
        for (int i = 0; i < certs.length; i++) {
            System.out.println(" certificate " + i);
            System.out.println("    " + certs[i]);
        }
    }

    if (certs != null && certs.length > 0) {

        // If the first certificate is an X509Certificate
        if (certs[0] instanceof X509Certificate) {
            try {
                // Get the first certificate which represents the user
                X509Certificate certX509 = (X509Certificate) certs[0];

                // Create a principal
                X500Principal principal = new X500Principal(certX509
                    .getIssuerDN()
                    .getName());
                principals.add(principal);

                if (debug) {
                    System.out.println(" Principal added: " + principal);
                }
                // Create the certification path object and add it to the
                // public credential set
                CertificateFactory factory = CertificateFactory
                    .getInstance("X.509");
                java.security.cert.CertPath certPath = factory
                    .generateCertPath(Arrays.asList(certs));
                publicCreds.add(certPath);

                // Add the private credential to the private credential set
                privateCreds.add(new X500PrivateCredential(certX509,
                    privateKey, userName));

            }
            catch (CertificateException ce) {
                SecurityException se = new SecurityException();
                se.initCause(ce);
                throw se;
            }
        }
        else {
            // The first certificate is not an X509Certificate
            // We just add the certificate to the public credential set
            // and the private key to the private credential set.
            publicCreds.add(certs[0]);
            privateCreds.add(privateKey);
        }
    }
}

```

```

    }
}
}

```

Using the LDAP authenticator plug-in

You are provided with the `com.ibm.websphere.objectgrid.security.plugins.builtins.LDAPAuthenticator` default implementation to handle the user name and password authentication to an LDAP server. This implementation uses the `LDAPLogin` login module to log the user into a Lightweight Directory Access Protocol (LDAP) server. The following snippet demonstrates how the `authenticate` method is implemented:

```

/**
 * @see com.ibm.ws.objectgrid.security.plugins.Authenticator#
 * authenticate(LDAPLogin)
 */
public Subject authenticate(Credential credential) throws
InvalidCredentialException, ExpiredCredentialException {

    UserPasswordCredential cred = (UserPasswordCredential) credential;
    LoginContext lc = null;
    try {
        lc = new LoginContext("LDAPLogin",
            new UserPasswordCallbackHandlerImpl(cred.getUserName(),
                cred.getPassword().toCharArray()));

        lc.login();

        Subject subject = lc.getSubject();

        return subject;
    }
    catch (LoginException le) {
        throw new InvalidCredentialException(le);
    }
    catch (IllegalArgumentException ile) {
        throw new InvalidCredentialException(ile);
    }
}

```

Also, eXtreme Scale ships a login module `com.ibm.websphere.objectgrid.security.plugins.builtins.LDAPLoginModule` for this purpose. You must provide the following two options in the JAAS login configuration file.

- `providerURL`: The LDAP server provider URL
- `factoryClass`: The LDAP context factory implementation class

The `LDAPLoginModule` module calls the `com.ibm.websphere.objectgrid.security.plugins.builtins.LDAPAuthenticationHelper.authenticate` method. The following code snippet shows how you can implement the `authenticate` method of the `LDAPAuthenticationHelper`.

```

/**
 * Authenticate the user to the LDAP directory.
 * @param user the user ID, e.g., uid=xxxxxx,c=us,ou=bluepages,o=ibm.com
 * @param pwd the password
 *
 * @throws NamingException
 */
public String[] authenticate(String user, String pwd)
throws NamingException {
    Hashtable env = new Hashtable();
    env.put(Context.INITIAL_CONTEXT_FACTORY, factoryClass);
    env.put(Context.PROVIDER_URL, providerURL);
    env.put(Context.SECURITY_PRINCIPAL, user);
    env.put(Context.SECURITY_CREDENTIALS, pwd);
    env.put(Context.SECURITY_AUTHENTICATION, "simple");
}

```

```

InitialContext initialContext = new InitialContext(env);

// Look up for the user
DirContext dirCtx = (DirContext) initialContext.lookup(user);

String uid = null;
int iComma = user.indexOf(",");
int iEqual = user.indexOf("=");
if (iComma > 0 && iComma > 0) {
    uid = user.substring(iEqual + 1, iComma);
}
else {
    uid = user;
}

Attributes attributes = dirCtx.getAttributes("");

// Check the UID
String thisUID = (String) (attributes.get(UID).get());

String thisDept = (String) (attributes.get(HR_DEPT).get());

if (thisUID.equals(uid)) {
    return new String[] { thisUID, thisDept };
}
else {
    return null;
}
}

```

If authentication succeeds, the ID and password are considered valid. Then the login module gets the ID information and department information from this authenticate method. The login module creates two principals: SimpleUserPrincipal and SimpleDeptPrincipal. You can use the authenticated subject for group authorization (in this case, the department is a group) and individual authorization.

The following example shows a login module configuration that is used to log in to the LDAP server:

```

LDAPLogin { com.ibm.websphere.objectgrid.security.plugins.builtins.LDAPLoginModule required
    providerURL="ldap://directory.acme.com:389/"
    factoryClass="com.sun.jndi.ldap.LdapCtxFactory";
};

```

In the previous configuration, the LDAP server points to the ldap://directory.acme.com:389/server. Change this setting to your LDAP server. This login module uses the provided ID and password to connect to the LDAP server. This implementation is for testing purposes only.

Using the WebSphere Application Server authenticator plug-in

Also, eXtreme Scale provides the com.ibm.websphere.objectgrid.security.plugins.builtins.WSTokenAuthenticator built-in implementation to use the WebSphere Application Server security infrastructure. This built-in implementation can be used when the following conditions are true.

1. WebSphere Application Server global security is turned on.
2. All eXtreme Scale clients and servers are launched in WebSphere Application Server JVMs.
3. These application servers are in the same security domain.

4. The eXtreme Scale client is already authenticated in WebSphere Application Server.

The client can use the `com.ibm.websphere.objectgrid.security.plugins.builtins.WSTokenCredentialGenerator` class to generate a credential. The server uses this Authenticator implementation class to authenticate the credential. If the token is authenticated successfully, a Subject object returns.

This scenario takes advantage of the fact that the client has already been authenticated. Because the application servers that have the servers are in the same security domain as the application servers that house the clients, the security tokens can be propagated from the client to the server so that the same user registry does not need to be authenticated again.

Using the Tivoli Access Manager authenticator plug-in

Tivoli Access Manager is used widely as a security server. You can also implement Authenticator using the Tivoli Access Manager's provided login modules.

To authenticate a user for Tivoli Access Manager, apply the the `com.tivoli.mts.PDLoginModule` login module, which requires that the calling application provide the following information:

1. A principal name, specified as either a short name or an X.500 name (DN)
2. A password

The login module authenticates the principal and returns the Tivoli Access Manager credential. The login module expects the calling application to provide the following information:

1. The user name, through a `javax.security.auth.callback.NameCallback` object.
2. The password, through a `javax.security.auth.callback.PasswordCallback` object.

When the Tivoli Access Manager credential is successfully retrieved, the JAAS LoginModule creates a Subject and a PDPrincipal. No built-in for Tivoli Access Manager authentication is provided, because it is just with the PDLoginModule module. See the IBM Tivoli Access Manager Authorization Java Classes Developer Reference for more details.

Connecting to WebSphere eXtreme Scale securely

To connect an eXtreme Scale client to a server securely, you can use any connect method in the ObjectGridManager interface which takes a ClientSecurityConfiguration object. The following is a brief example.

```
public ClientClusterContext connect(String catalogServerAddresses,  
    ClientSecurityConfiguration securityProps,  
    URL overrideObjectGridXml) throws ConnectException;
```

This method takes a parameter of the ClientSecurityConfiguration type, which is an interface representing a client security configuration. You can use `com.ibm.websphere.objectgrid.security.config.ClientSecurityConfigurationFactory` public API to create an instance with default values, or you can create an instance by passing the WebSphere eXtreme Scale client property file. This file contains the following properties that are related to authentication. The value marked with a plus sign (+) is the default.

- `securityEnabled (true, false+)`: This property indicates if security is enabled. When a client connects to a server, the `securityEnabled` value on the client and

server side must be both true or both false. For example, if the connected server security is enabled, the client has to set this property to true to connect to the server.

- `authenticationRetryCount` (an integer value, 0+): This property determines how many retries are attempted for login when a credential is expired. If the value is 0, no retries are attempted. The authentication retry only applies to the case when the credential is expired. If the credential is not valid, there is no retry. Your application is responsible for retrying the operation.

After you create a `com.ibm.websphere.objectgrid.security.config.ClientSecurityConfiguration` object, set the `credentialGenerator` object on the client using the following method:

```
/**
 * Set the {@link CredentialGenerator} object for this client.
 * @param generator the CredentialGenerator object associated with this client
 */
void setCredentialGenerator(CredentialGenerator generator);
```

You can set the `CredentialGenerator` object in the WebSphere eXtreme Scale client property file too, as follows.

- `credentialGeneratorClass`: The class implementation name for the `CredentialGenerator` object. It must have a default constructor.
- `credentialGeneratorProps`: The properties for the `CredentialGenerator` class. If the value is not null, it is set to the constructed `CredentialGenerator` object using the `setProperties(String)` method.

Here is a sample to instantiate a `ClientSecurityConfiguration` and then use it to connect to the server.

```
/**
 * Get a secure ClientClusterContext
 * @return a secure ClientClusterContext object
 */
protected ClientClusterContext connect() throws ConnectException {
    ClientSecurityConfiguration csConfig = ClientSecurityConfigurationFactory
        .getClientSecurityConfiguration("/properties/security.ogclient.props");

    UserPasswordCredentialGenerator gen= new
        UserPasswordCredentialGenerator("manager", "manager1");

    csConfig.setCredentialGenerator(gen);

    return objectGridManager.connect(csConfig, null);
}
```

When the `connect` is called, the WebSphere eXtreme Scale client calls the `CredentialGenerator.getCredential` method to get the client credential. This credential is sent along with the connect request to the server for authentication.

Using a different CredentialGenerator instance per session

In some cases, a WebSphere eXtreme Scale client represents just one client identity, but in others, it might represent multiple identities. Here is one scenario for the latter case: An WebSphere eXtreme Scale client is created and shared in a Web server. All servlets in this Web server use this one WebSphere eXtreme Scale client. Because every servlet represents a different Web client, use different credentials when sending requests to WebSphere eXtreme Scale servers.

WebSphere eXtreme Scale provides for changing the credential on the session level. Every session can use a different CredentialGenerator object. Therefore, the previous scenarios can be implemented by letting the servlet get a session with a different CredentialGenerator object. The following example illustrates the ObjectGrid.getSession(CredentialGenerator) method in the ObjectGridManager interface.

```
/**
 * Get a session using a <code>CredentialGenerator</code>.
 * <p>
 * This method can only be called by the ObjectGrid client in an ObjectGrid
 * client server environment. If ObjectGrid is used in a local model, that is,
 * within the same JVM with no client or server existing, <code>getSession(Subject)</code>
 * or the <code>SubjectSource</code> plugin should be used to secure the ObjectGrid.
 *
 * <p>If the <code>initialize()</code> method has not been invoked prior to
 * the first <code>getSession</code> invocation, an implicit initialization
 * will occur. This ensures that all of the configuration is complete
 * before any runtime usage is required.</p>
 *
 * @param credGen A <code>CredentialGenerator</code> for generating a credential
 *               for the session returned.
 *
 * @return An instance of <code>Session</code>
 *
 * @throws ObjectGridException if an error occurs during processing
 * @throws TransactionCallbackException if the <code>TransactionCallback</code>
 *         throws an exception
 * @throws IllegalStateException if this method is called after the
 *         <code>destroy()</code> method is called.
 *
 * @see #destroy()
 * @see #initialize()
 * @see CredentialGenerator
 * @see Session
 * @since WAS XD 6.0.1
 */
Session getSession(CredentialGenerator credGen) throws
ObjectGridException, TransactionCallbackException;
```

The following is an example:

```
ObjectGridManager ogManager = ObjectGridManagerFactory.getObjectGridManager();
CredentialGenerator credGenManager = new UserPasswordCredentialGenerator("manager", "xxxxxx");
CredentialGenerator credGenEmployee = new UserPasswordCredentialGenerator("employee", "xxxxxx");

ObjectGrid og = ogManager.getObjectGrid(ctx, "accounting");

// Get a session with CredentialGenerator;
Session session = og.getSession(credGenManager );

// Get the employee map
ObjectMap om = session.getMap("employee");

// start a transaction.
session.begin();

Object rec1 = map.get("xxxxxx");

session.commit();

// Get another session with a different CredentialGenerator;
session = og.getSession(credGenEmployee );

// Get the employee map
om = session.getMap("employee");

// start a transaction.
session.begin();

Object rec2 = map.get("xxxxxx");

session.commit();
```

If you use the ObjectGrid.getSession method to get a Session object, the session uses the CredentialGenerator object set on the ClientConfigurationSecurity object. The ObjectGrid.getSession(CredentialGenerator) method overrides the CredentialGenerator set in the ClientSecurityConfiguration object.

If you can reuse the Session object, a performance gain results. However, calling the ObjectGrid.getSession(CredentialGenerator) method is not very expensive. The

major overhead is the increased object garbage collection time. Make sure that you release the references after you are done with the Session objects. Generally, if your Session object can share the identity, try to reuse the Session object. If not, use the `ObjectGrid.getSession(CredentialGenerator)` method.

Client authorization programming

WebSphere eXtreme Scale supports Java Authentication and Authorization Service (JAAS) authorization out-of-the-box and also supports custom authorization using the `ObjectGridAuthorization` interface.

The `ObjectGridAuthorization` plug-in is used to authorize `ObjectGrid`, `ObjectMap` and `JavaMap` accesses to the Principals represented by a Subject object in a custom way. A typical implementation of this plug-in is to retrieve the Principals from the Subject object, and then check whether or not the specified permissions are granted to the Principals.

A permission passed to the `checkPermission(Subject, Permission)` method can be one of the following permissions:

- `MapPermission`
- `ObjectGridPermission`
- `ServerMapPermission`
- `AgentPermission`

Refer to `ObjectGridAuthorization` API documentation for more details.

MapPermission

The `com.ibm.websphere.objectgrid.security.MapPermission` public class represents permissions to the `ObjectGrid` resources, specifically the methods of `ObjectMap` or `JavaMap` interfaces. WebSphere eXtreme Scale defines the following permission strings to access the methods of `ObjectMap` and `JavaMap`:

- **read**: Permission to read the data from the map. The integer constant is defined as `MapPermission.READ`.
- **write**: Permission to update the data in the map. The integer constant is defined as `MapPermission.WRITE`.
- **insert**: Permission to insert the data into the map. The integer constant is defined as `MapPermission.INSERT`.
- **remove**: Permission to remove the data from the map. The integer constant is defined as `MapPermission.REMOVE`.
- **invalidate**: Permission to invalidate the data from the map. The integer constant is defined as `MapPermission.INVALIDATE`.
- **all**: All above permissions: read, write, insert, remote, and invalidate. The integer constant is defined as `MapPermission.ALL`.

Refer to `MapPermission` API documentation for more details.

You can construct a `MapPermission` object by passing the fully qualified `ObjectGrid` map name (in format `[ObjectGrid_name].[ObjectMap_name]`) and the permission string or integer value. A permission string can be a comma-delimited string of the previous permission strings such as read, insert, or it can be all. A permission integer value can be any previously mentioned permission integer

constants or a mathematical value of several integer permission constants, such as `MapPermission.READ` | `MapPermission.WRITE`.

The authorization occurs when an `ObjectMap` or `JavaMap` method is called. The eXtreme Scale runtime checks different permissions for different methods. If the required permissions are not granted to the client, an `AccessControlException` results.

Table 6. List of methods and the required MapPermission

| Permission | ObjectMap/JavaMap |
|------------|--|
| read | boolean <code>containsKey(Object)</code> |
| | boolean <code>equals(Object)</code> |
| | Object <code>get(Object)</code> |
| | Object <code>get(Object, Serializable)</code> |
| | List <code>getAll(List)</code> |
| | List <code>getAll(List keyList, Serializable)</code> |
| | List <code>getAllForUpdate(List)</code> |
| | List <code>getAllForUpdate(List, Serializable)</code> |
| | Object <code>getForUpdate(Object)</code> |
| | Object <code>getForUpdate(Object, Serializable)</code> |
| | public Object <code>getNextKey(long)</code> |
| write | Object <code>put(Object key, Object value)</code> |
| | void <code>put(Object, Object, Serializable)</code> |
| | void <code>putAll(Map)</code> |
| | void <code>putAll(Map, Serializable)</code> |
| | void <code>update(Object, Object)</code> |
| | void <code>update(Object, Object, Serializable)</code> |
| insert | public void <code>insert (Object, Object)</code> |
| | void <code>insert(Object, Object, Serializable)</code> |
| remove | Object <code>remove (Object)</code> |
| | void <code>removeAll(Collection)</code> |
| | void <code>clear()</code> |
| invalidate | public void <code>invalidate (Object, boolean)</code> |
| | void <code>invalidateAll(Collection, boolean)</code> |
| | void <code>invalidateUsingKeyword(Serializable)</code> |
| | int <code>setTimeToLive(int)</code> |

Authorization is based solely on which method is used, rather than what the method really does. For example, a `put` method can insert or update a record based on whether the record exists. However, the insert or update cases are not distinguished.

An operation type can be achieved by combinations of other types. For example, an update can be achieved by a remove and then an insert. Consider these combinations when designing your authorization policies.

ObjectGridPermission

A `com.ibm.websphere.objectgrid.security.ObjectGridPermission` represents permissions to the ObjectGrid:

- Query: permission to create an object query or entity query. The integer constant is defined as `ObjectGridPermission.QUERY`.
- Dynamic map: permission to create a dynamic map based on the map template. The integer constant is defined as `ObjectGridPermission.DYNAMIC_MAP`.

Refer to ObjectGridPermission API documentation for more details.

The following table summarizes the methods and the required ObjectGridPermission:

Table 7. List of methods and the required ObjectGridPermission

| Permission action | Methods |
|-------------------|--|
| query | <code>com.ibm.websphere.objectgrid.Session.createObjectQuery(String)</code> |
| query | <code>com.ibm.websphere.objectgrid.em.EntityManager.createQuery(String)</code> |
| dynamicmap | <code>com.ibm.websphere.objectgrid.Session.getMap(String)</code> |

ServerMapPermission

An `ServerMapPermission` represents permissions to an ObjectMap hosted in a server. The name of the permission is the full name of the ObjectGrid map name. It has the following actions:

- replicate: permission to replicate a server map to near cache
- dynamicIndex: permission for a client to create or remove a dynamic index on a server

Refer to ServerMapPermission API documentation for more details. The detailed methods, which require different ServerMapPermission, are listed in the following table:

Table 8. Permissions to a server-hosted ObjectMap

| Permission action | Methods |
|-------------------|--|
| replicate | <code>com.ibm.websphere.objectgrid.ClientReplicableMap.enableClientReplication(Mode, int[], ReplicationMapListener)</code> |
| dynamicIndex | <code>com.ibm.websphere.objectgrid.BackingMap.createDynamicIndex(String, boolean, String, DynamicIndexCallback)</code> |
| dynamicIndex | <code>com.ibm.websphere.objectgrid.BackingMap.removeDynamicIndex(String)</code> |

AgentPermission

An `AgentPermission` represents permissions to the datagrid agents. The name of the permission is the full name of the ObjectGrid map, and the action is a comma-delimited string of agent implementation class names or package names.

Refer to AgentPermission API documentation for more information.

The following methods in the class `com.ibm.websphere.objectgrid.datagrid.AgentManager` require AgentPermission.

```
com.ibm.websphere.objectgrid.datagrid.AgentManager#callMapAgent(MapGridAgent, Collection)
com.ibm.websphere.objectgrid.datagrid.AgentManager#callMapAgent(MapGridAgent)
com.ibm.websphere.objectgrid.datagrid.AgentManager#callReduceAgent(ReduceGridAgent, Collection)
com.ibm.websphere.objectgrid.datagrid.AgentManager#callReduceAgent(ReduceGridAgent, Collection)
```

Authorization mechanisms

WebSphere eXtreme Scale supports two kinds of authorization mechanisms: Java Authentication and Authorization Service (JAAS) authorization and custom authorization. These mechanisms apply to all authorizations. JAAS authorization augments the Java security policies with user-centric access controls. Permissions can be granted based not just on what code is running, but also on who is running it. JAAS authorization is part of the SDK Version 1.4 and later.

Additionally, WebSphere eXtreme Scale also supports custom authorization with the following plug-in:

- ObjectGridAuthorization: custom way to authorize access to all artifacts.

You can implement your own authorization mechanism if you do not want to use JAAS authorization. By using a custom authorization mechanism, you can use the policy database, policy server, or Tivoli Access Manager to manage the authorizations.

You can configure the authorization mechanism in two ways:

- XML configuration

You can use the ObjectGrid XML file to define an ObjectGrid and set the authorization mechanism to either AUTHORIZATION_MECHANISM_JAAS or AUTHORIZATION_MECHANISM_CUSTOM. Here is the secure-objectgrid-definition.xml file that is used in the enterprise application ObjectGridSample:

```
<objectGrids>
  <objectGrid name="secureClusterObjectGrid" securityEnabled="true"
    authorizationMechanism="AUTHORIZATION_MECHANISM_JAAS">
    <bean id="TransactionCallback"
      classname="com.ibm.websphere.samples.objectgrid.HeapTransactionCallback" />
    ...
  </objectGrids>
```

- Programmatic configuration

If you want to create an ObjectGrid using method `ObjectGrid.setAuthorizationMechanism(int)`, you can call the following method to set the authorization mechanism. Calling this method applies only to the local WebSphere eXtreme Scale programming model when you directly instantiate the ObjectGrid instance:

```
/**
 * Set the authorization Mechanism. The default is
 * com.ibm.websphere.objectgrid.security.SecurityConstants.
 * AUTHORIZATION_MECHANISM_JAAS.
 * @param authMechanism the map authorization mechanism
 */
void setAuthorizationMechanism(int authMechanism);
```

JAAS authorization

A `javax.security.auth.Subject` object represents an authenticated user. A Subject is comprised of a set of principals, and each Principal represents an identity for that user. For example, a Subject can have a name principal, for example, Joe Smith, and a group principal, for example, manager.

Using the JAAS authorization policy, permissions can be granted to specific Principals. WebSphere eXtreme Scale associates the Subject with the current access control context. For each call to the `ObjectMap` or `Javamap` method, the Java runtime automatically determines if the policy grants the required permission only

to a specific Principal and if so, the operation is allowed only if the Subject associated with the access control context contains the designated Principal.

You must be familiar with the policy syntax of the policy file. For detailed description of JAAS authorization, refer to the JAAS Reference Guide.

WebSphere eXtreme Scale has a special code base that is used for checking the JAAS authorization to the ObjectMap and JavaMap method calls. This special code base is `http://www.ibm.com/com/ibm/ws/objectgrid/security/PrivilegedAction`. Use this code base when granting ObjectMap or JavaMap permissions to principals. This special code was created because the Java archive (JAR) file for eXtreme Scale is granted with all permissions.

The template of the policy to grant the MapPermission permission is:

```
grant codeBase "http://www.ibm.com/com/ibm/ws/objectgrid/security/PrivilegedAction"
  <Principal field(s)>{
    permission com.ibm.websphere.objectgrid.security.MapPermission
      "[ObjectGrid_name].[ObjectMap_name]", "action";
    ....
    permission com.ibm.websphere.objectgrid.security.MapPermission
      "[ObjectGrid_name].[ObjectMap_name]", "action";
  };
```

A Principal field looks like the following example:

```
principal Principal_class "principal_name"
```

In this policy, only insert and read permissions are granted to these four maps to a certain principal. The other policy file, `fullAccessAuth.policy`, grants all permissions to these maps to a principal. Before running the application, change the `principal_name` and `principal class` to appropriate values. The value of the `principal_name` depends on the user registry. For example, if local OS is used as user registry, the machine name is `MACH1`, the user ID is `user1`, and the `principal_name` is `MACH1/user1`.

The JAAS authorization policy can be put directly into the Java policy file, or it can be put in a separate JAAS authorization file and then set in either of two ways:

- Use the following JVM argument:
`-Djava.security.auth.policy=file:[JAAS_AUTH_POLICY_FILE]`
- Use the following property in the `java.security` file:
`-Dauth.policy.url.x=file:[JAAS_AUTH_POLICY_FILE]`

Custom ObjectGrid authorization

ObjectGridAuthorization plug-in is used to authorize ObjectGrid, ObjectMap and JavaMap accesses to the Principals represented by a Subject object in a custom way. A typical implementation of this plug-in is to retrieve the Principals from the Subject object, and then check whether or not the specified permissions are granted to the Principals.

A permission passed to the `checkPermission(Subject, Permission)` method could be one of the following:

- MapPermission
- ObjectGridPermission
- AgentPermission
- ServerMapPermission

Refer to ObjectGridAuthorization API documentation for more details.

The ObjectGridAuthorization plug-in can be configured in the following ways:

- XML configuration

You can use the ObjectGrid XML file to define an ObjectAuthorization plug-in. Here is an example:

```
<objectGrids>
  <objectGrid name="secureClusterObjectGrid" securityEnabled="true"
    authorizationMechanism="AUTHORIZATION_MECHANISM_CUSTOM">
    ...
  <bean id="ObjectGridAuthorization"
    className="com.acme.ObjectGridAuthorizationImpl" />
</objectGrids>
```

- Programmatic configuration

If you want to create an ObjectGrid using the API method `ObjectGrid.setObjectGridAuthorization(ObjectGridAuthorization)`, you can call the following method to set the authorization plug-in. This method only applies to the local eXtreme Scale programming model when you directly instantiate the ObjectGrid instance.

```
/**
 * Sets the <code>ObjectGridAuthorization</code> for this ObjectGrid instance.
 * <p>
 * Passing <code>null</code> to this method removes a previously set
 * <code>ObjectGridAuthorization</code> object from an earlier invocation of this method
 * and indicates that this <code>ObjectGrid</code> is not associated with a
 * <code>ObjectGridAuthorization</code> object.
 * <p>
 * This method should only be used when ObjectGrid security is enabled. If
 * the ObjectGrid security is disabled, the provided <code>ObjectGridAuthorization</code> object
 * will not be used.
 * <p>
 * A <code>ObjectGridAuthorization</code> plugin can be used to authorize
 * access to the ObjectGrid and maps. Please refer to <code>ObjectGridAuthorization</code> for more details.
 * <p>
 * As of XD 6.1, the <code>setMapAuthorization</code> is deprecated and
 * <code>setObjectGridAuthorization</code> is recommended for use. However,
 * if both <code>MapAuthorization</code> plugin and <code>ObjectGridAuthorization</code> plugin
 * are used, ObjectGrid will use the provided <code>MapAuthorization</code> to authorize map accesses,
 * even though it is deprecated.
 * <p>
 * Note, to avoid an <code>IllegalStateException</code>, this method must be
 * called prior to the <code>initialize</code> method. Also, keep in mind
 * that the <code>getSession</code> methods implicitly call the
 * <code>initialize</code> method if it has yet to be called by the
 * application.
 * @param ogAuthorization the <code>ObjectGridAuthorization</code> plugin
 * @throws IllegalStateException if this method is called after the
 * <code>initialize</code> method is called.
 * @see #initialize()
 * @see ObjectGridAuthorization
 * @since WAS XD 6.1
 */
void setObjectGridAuthorization(ObjectGridAuthorization ogAuthorization);
```

Implementing ObjectGridAuthorization

The boolean `checkPermission(Subject subject, Permission permission)` method of the ObjectGridAuthorization interface is called by the WebSphere eXtreme Scale run time to check whether the passed-in subject object has the passed-in permission. The implementation of the ObjectGridAuthorization interface returns true if the object has the permission, and false if not.

A typical implementation of this plug-in is to retrieve the principals from the Subject object and check whether the specified permissions are granted to the principals by consulting specific policies. These policies are defined by users. For example, the policies can be defined in a database, a plain file, or a Tivoli Access Manager policy server.

For example, we can use Tivoli Access Manager policy server to manage the authorization policy and use its API to authorize the access. For how to use Tivoli Access Manager Authorization APIs, refer to the IBM Tivoli Access Manager Authorization Java Classes Developer Reference for more details.

This sample implementation has the following assumptions:

- Only check authorization for MapPermission. For other permissions, always return true.
- The Subject object contains a com.tivoli.mts.PDPrincipal principal.
- The Tivoli Access Manager policy server has defined the following permissions for the ObjectMap or JavaMap name object. The object that is defined in the policy server must have the same name as the ObjectMap or JavaMap name in the format of [ObjectGrid_name].[ObjectMap_name]. The permission is the first character of the permission strings that are defined in the MapPermission permission. For example, the permission "r" that is defined in the policy server represents the read permission to the ObjectMap map.

The following code snippet demonstrates how to implement the checkPermission method:

```
/**
 * @see com.ibm.websphere.objectgrid.security.plugins.
 * MapAuthorization#checkPermission
 * (javax.security.auth.Subject, com.ibm.websphere.objectgrid.security.
 * MapPermission)
 */
public boolean checkPermission(final Subject subject,
    Permission p) {

    // For non-MapPermission, we always authorize.
    if (!(p instanceof MapPermission)){
        return true;
    }

    MapPermission permission = (MapPermission) p;

    String[] str = permission.getParsedNames();

    StringBuffer pdPermissionStr = new StringBuffer(5);
    for (int i=0; i<str.length; i++) {
        pdPermissionStr.append(str[i].substring(0,1));
    }

    PDPermission pdPerm = new PDPermission(permission.getName(),
        pdPermissionStr.toString());

    Set principals = subject.getPrincipals();

    Iterator iter= principals.iterator();
    while(iter.hasNext()) {
        try {
            PDPrincipal principal = (PDPrincipal) iter.next();
            if (principal.implies(pdPerm)) {
                return true;
            }
        }
        catch (ClassCastException cce) {
            // Handle exception
        }
    }
    return false;
}
```

Data grid authentication

You can use the secure token manager plug-in to enable server-to-server authentication, which requires you to implement the `SecureTokenManager` interface.

The `generateToken(Object)` method takes an object to protect, and then generates a token that cannot be understood by others. The `verifyTokens(byte[])` method does the reverse process: it converts the token back to the original object.

A simple `SecureTokenManager` implementation uses a simple encoding algorithm, such as a XOR algorithm, to encode the object in serialized form and then use corresponding decoding algorithm to decode the token. This implementation is not secure and is easy to break.

WebSphere eXtreme Scale default implementation

WebSphere eXtreme Scale provides an immediately available implementation for this interface. This default implementation uses a key pair to sign and verify the signature, and uses a secret key to encrypt the content. Every server has a JCKES type keystore to store the key pair, a private key and public key, and a secret key. The keystore has to be the JCKES type to store secret keys. These keys are used to encrypt and sign or verify the secret string on the sending end. Also, the token is associated with an expiration time. On the receiving end, the data is verified, decrypted, and compared to the receiver secret string. Secure Sockets Layer (SSL) communication protocols are not required between a pair of servers for authentication because the private keys and public keys serve the same purpose. However, if server communication is not encrypted, the data can be stolen by looking at the communication. Because the token expires soon, the replay attack threat is minimized. This possibility is significantly decreased if all servers are deployed behind a firewall.

The disadvantage of this approach is that the WebSphere eXtreme Scale administrators have to generate keys and transport them to all servers, which can cause security breach during transportation.

Local security

WebSphere eXtreme Scale provides several security endpoints to allow you to integrate custom mechanisms. In the local programming model, the main security function is authorization, and has no authentication support. You must authenticate outside of WebSphere Application Server. However, there are provided plug-ins to obtain and validate Subject objects.

Enabling security

The following list provides the two ways in which local security is enabled:

- **XML Configuration** You can use the ObjectGrid XML file to define an ObjectGrid and enable the security for that ObjectGrid. The following file is the `secure-objectgrid-definition.xml` file that is used in the ObjectGridSample enterprise application sample. In this XML file, security is enabled by setting the `securityEnabled` attribute to true.

```

<objectGrids>
  <objectGrid name="secureClusterObjectGrid" securityEnabled="true"
    authorizationMechanism="AUTHORIZATION_MECHANISM_JASS">
    ...
</objectGrids>

```

- **Programming** If you want to create an ObjectGrid using the API method ObjectGrid.setSecurityEnabled(), call the following method on the ObjectGrid interface to enable security.

```

/**
 * Enable the ObjectGrid security
 */
void setSecurityEnabled();

```

Authentication

In the local programming model, eXtreme Scale does not provide any authentication mechanism, but relies on the environment, either application servers or applications, for authentication. When eXtreme Scale is used in WebSphere Application Server or WebSphere Extended Deployment, applications can use the WebSphere Application Server security authentication mechanism. When eXtreme Scale is running in a Java 2 Platform, Standard Edition (J2SE) environment, the application has to manage authentications with Java Authentication and Authorization Service (JAAS) authentication or other authentication mechanisms. For more information about using JAAS authentication, see the JAAS reference guide. The contract between an application and an ObjectGrid instance is the javax.security.auth.Subject object. After the client is authenticated by the application server or the application, the application can retrieve the authenticated javax.security.auth.Subject object and use this Subject object to get a session from the ObjectGrid instance by calling the ObjectGrid.getSession(Subject) method. This Subject object is used to authorize accesses to the map data. This contract is called a subject passing mechanism. The following example illustrates the ObjectGrid.getSession(Subject) API.

```

/**
 * This API allows the cache to use a specific subject rather than the one
 * configured on the ObjectGrid to get a session.
 * @param subject
 * @return An instance of Session
 * @throws ObjectGridException
 * @throws TransactionCallbackException
 * @throws InvalidSubjectException the subject passed in is not valid based
 * on the SubjectValidation mechanism.
 */
public Session getSession(Subject subject)
throws ObjectGridException, TransactionCallbackException, InvalidSubjectException;

```

The ObjectGrid.getSession() method in the ObjectGrid interface can also be used to get a Session object:

```

/**
 * This method returns a Session object that can be used by a single thread at a time.
 * You cannot share this Session object between threads without placing a
 * critical section around it. While the core framework allows the object to move
 * between threads, the TransactionCallback and Loader might prevent this usage,
 * especially in J2EE environments. When security is enabled, this method uses the
 * SubjectSource to get a Subject object.
 *
 * If the initialize method has not been invoked prior to the first
 * getSession invocation, then an implicit initialization occurs. This
 * initialization ensures that all of the configuration is complete before
 * any runtime usage is required.
 *
 * @see #initialize()
 * @return An instance of Session
 * @throws ObjectGridException

```



```

* @throws TransactionCallbackException
* @throws IllegalStateException if this method is called after the
*       destroy() method is called.
*/
public Session getSession()
throws ObjectGridException, TransactionCallbackException;

```

As the API documentation specifies, when security is enabled, this method uses the SubjectSource plug-in to get a Subject object. The SubjectSource plug-in is one of the security plug-ins defined in eXtreme Scale to support propagating Subject objects. See Security-related plug-ins for more information. The getSession(Subject) method can be called on the local ObjectGrid instance only. If you call the getSession(Subject) method on a client side in a distributed eXtreme Scale configuration, an IllegalStateException results.

Security plug-ins

WebSphere eXtreme Scale provides two security plug-ins that are related to the subject passing mechanism: the SubjectSource and SubjectValidation plug-ins.

SubjectSource plug-in

The SubjectSource plug-in, represented by the com.ibm.websphere.objectgrid.security.plugins.SubjectSource interface, is a plug-in that is used to get a Subject object from an eXtreme Scale running environment. This environment can be an application using the ObjectGrid or an application server that hosts the application. Consider the SubjectSource plug-in an alternative to the subject passing mechanism. Using the subject passing mechanism, the application retrieves the Subject object and uses it to get the ObjectGrid session object. With the SubjectSource plug-in, the eXtreme Scale runtime retrieves the Subject object and uses it to get the session object. The subject passing mechanism gives the control of Subject objects to applications, while the SubjectSource plug-in mechanism frees applications from retrieving the Subject object. You can use the SubjectSource plug-in to get a Subject object that represents an eXtreme Scale client that is used for authorization. When the ObjectGrid.getSession method is called, the Subject getSubject throws an ObjectGridSecurityException if security is enabled. WebSphere eXtreme Scale provides a default implementation of this plug-in: com.ibm.websphere.objectgrid.security.plugins.builtins.WSSubjectSourceImpl. This implementation can be used to retrieve a caller subject or a RunAs subject from the thread when an application is running in WebSphere Application Server. You can configure this class in your ObjectGrid descriptor XML file as the SubjectSource implementation class when using eXtreme Scale in WebSphere Application Server. The following code snippet shows the main flow of the WSSubjectSourceImpl.getSubject method.

```

Subject s = null;
try {
    if (finalType == RUN_AS_SUBJECT) {
        // get the RunAs subject
        s = com.ibm.websphere.security.auth.WSSubject.getRunAsSubject();
    }
    else if (finalType == CALLER_SUBJECT) {
        // get the callersubject
        s = com.ibm.websphere.security.auth.WSSubject.getCallerSubject();
    }
}
catch (WSSecurityException wse) {
    throw new ObjectGridSecurityException(wse);
}

return s;

```

For other details, refer to the API documentation for the SubjectSource plug-in and the WSSubjectSourceImpl implementation.

SubjectValidation plug-in

The SubjectValidation plug-in, which is represented by the `com.ibm.websphere.objectgrid.security.plugins.SubjectValidation` interface, is another security plug-in. The SubjectValidation plug-in can be used to validate that a `javax.security.auth.Subject`, either passed to the ObjectGrid or retrieved by the SubjectSource plug-in, is a valid Subject that has not been tampered with.

The `SubjectValidation.validateSubject(Subject)` method in the SubjectValidation interface takes a Subject object and returns a Subject object. Whether a Subject object is considered valid and which Subject object is returned are all up to your implementations. If the Subject object is not valid, an `InvalidSubjectException` results.

You can use this plug-in if you do not trust the Subject object that is passed to this method. This case is rare considering that you trust the application developers who develop the code to retrieve the Subject object.

An implementation of this plug-in needs support from the Subject object creator because only the creator knows if the Subject object has been tampered with. However, some subject creator might not know if the Subject has been tampered with. In this case, this plug-in is not useful.

WebSphere eXtreme Scale provides a default implementation of SubjectValidation: `com.ibm.websphere.objectgrid.security.plugins.builtins.WSSubjectValidationImpl`. You can use this implementation to validate the WebSphere Application Server-authenticated subject. You can configure this class as the SubjectValidation implementation class when using eXtreme Scale in WebSphere Application Server. The `WSSubjectValidationImpl` implementation considers a Subject object valid only if the credential token that is associated with this Subject has not been tampered with. You can change other parts of the Subject object. The `WSSubjectValidationImpl` implementation asks WebSphere Application Server for the original Subject corresponding to the credential token and returns the original Subject object as the validated Subject object. Therefore, the changes made to the Subject contents other than the credential token have no effects. The following code snippet shows the basic flow of the `WSSubjectValidationImpl.validateSubject(Subject)`.

```
// Create a LoginContext with scheme WSLogin and
// pass a Callback handler.
LoginContext lc = new LoginContext("WSLogin",
new WSCredTokenCallbackHandlerImpl(subject));

// When this method is called, the callback handler methods
// will be called to log the user in.
lc.login();

// Get the subject from the LoginContext
return lc.getSubject();
```

In the previous code snippet, a credential token callback handler object, `WSCredTokenCallbackHandlerImpl`, is created with the Subject object to validate. Then a `LoginContext` object is created with the login scheme `WSLogin`. When the

login method is called, WebSphere Application Server security retrieves the credential token from the Subject object and then returns the correspondent Subject as the validated Subject object.

For other details, refer to the Java APIs of SubjectValidation and WSSubjectValidationImpl implementation.

Plug-in configuration

You can configure the SubjectValidation plug-in and SubjectSource plug-in in two ways:

- **XML Configuration** You can use the ObjectGrid XML file to define an ObjectGrid and set these two plug-ins. Here is an example, in which the WSSubjectSourceImpl class is configured as the SubjectSource plug-in and the WSSubjectValidation class is configured as the SubjectValidation plug-in.

```
<objectGrids>
  <objectGrid name="secureClusterObjectGrid" securityEnabled="true"
    authorizationMechanism="AUTHORIZATION_MECHANISM_JAAS">
    <bean id="SubjectSource"
      className="com.ibm.websphere.objectgrid.security.plugins.builtins.
        WSSubjectSourceImpl" />
    <bean id="SubjectValidation"
      className="com.ibm.websphere.objectgrid.security.plugins.builtins.
        WSSubjectValidationImpl" />
    <bean id="TransactionCallback"
      className="com.ibm.websphere.samples.objectgrid.
        HeapTransactionCallback" />
    ...
  </objectGrids>
```

- **Programming** If you want to create an ObjectGrid through APIs, you can call the following methods to set the SubjectSource or SubjectValidation plug-ins.

```
**
 * Set the SubjectValidation plug-in for this ObjectGrid instance. A
 * SubjectValidation plug-in can be used to validate the Subject object
 * passed in as a valid Subject. Refer to {@link SubjectValidation}
 * for more details.
 * @param subjectValidation the SubjectValidation plug-in
 */
void setSubjectValidation(SubjectValidation subjectValidation);

/**
 * Set the SubjectSource plug-in. A SubjectSource plug-in can be used
 * to get a Subject object from the environment to represent the
 * ObjectGrid client.
 *
 * @param source the SubjectSource plug-in
 */
void setSubjectSource(SubjectSource source);
```

Write your own JAAS authentication code

You can write your own Java Authentication and Authorization Service (JAAS) authentication code to handle the authentication. You need to write your own login modules and then configure the login modules for your authentication module.

The login module receives information about a user and authenticates the user. This information can be anything that can identify the user. For example, the information can be a user ID and password, client certificate, and so on. After receiving the information, the login module verifies that the information represents

a valid subject and then creates a Subject object. Currently, several implementations of login modules are available to the public.

After a login module is written, configure this login module for the run time to use. You must configure a JAAS login module. This login module contains the login module and its authentication scheme. For example:

```
FileLogin
{
    com.acme.auth.FileLoginModule required
};
```

The authentication scheme is FileLogin and the login module is com.acme.auth.FileLoginModule. The required token indicates that the FileLoginModule module must validate this login or the entire scheme fails.

Setting the JAAS login module configuration file can be done in one of the following ways:

- Set the JAAS login module configuration file in the login.config.url property in the java.security file, for example:
login.config.url.1=file:\${java.home}/lib/security/file.login
- Set the JAAS login module configuration file from the command line by using the **-Djava.security.auth.login.config** Java virtual machine (JVM) arguments, for example, **-Djava.security.auth.login.config ==\$JAVA_HOME/lib/security/file.login**

If your code is running in WebSphere Application Server, you must configure the JAAS login in the administrative console and store this login configuration in the application server configuration. See Login configuration for Java Authentication and Authorization Service for details.

Chapter 9. Performance considerations for application developers

To improve performance for your in-memory data grid or database processing space, you can investigate several considerations such as tuning your Java virtual machine settings and using the best practices for product features such as locking, serialization, and query performance.

JVM tuning for WebSphere eXtreme Scale

You must take into account several specific aspects of Java virtual machine (JVM) tuning for WebSphere eXtreme Scale best performance. In most cases, few or no special JVM settings are required. If many objects are being stored in the data grid, adjust the heap size to an appropriate level to avoid running out of memory.

Tested platforms

Performance testing occurred primarily on AIX® (32 way), Linux (four way), and Windows (eight way) computers. With high-end AIX computers, you can test heavily multi-threaded scenarios to identify and fix contention points.

Garbage collection

WebSphere eXtreme Scale creates temporary objects that are associated with each transaction, such as request and response, and log sequence. Because these objects affect garbage collection efficiency, tuning garbage collection is critical.

For the IBM virtual machine for Java, use the **optavgpause** collector for high update rate scenarios (100% of transactions modify entries). The **gencon** collector works much better than the **optavgpause** collector for scenarios where data is updated relatively infrequently (10% of the time or less). Experiment with both collectors to see what works best in your scenario. Run with verbose garbage collection turned on to check the percentage of the time that is being spent collecting garbage. Scenarios have occurred where 80% of the time is spent in garbage collection until tuning fixed the problem.

All modern JVMs today use parallel garbage collection algorithms, which means that using more cores can reduce pauses in garbage collection. A physical server with eight cores has a faster garbage collection than a physical with four cores.

When the application must manage a large amount of data for each partition, then garbage collection might be a factor. A read mostly scenario performs even with large heaps (20 GB or more) if a generational collector is used. However, after the tenure heap fills, a pause proportional to the live heap size and the number of processors on the computer occurs. This pause can be large on smaller computers with large heaps.

Attention: If you are using a Sun JVM, adjustments to the default garbage collection and tuning policy might be necessary.

WebSphere eXtreme Scale supports WebSphere Real Time Java. With WebSphere Real Time Java, the transaction processing response for WebSphere eXtreme Scale is more consistent and predictable. As a result, the impact of garbage collection and

thread scheduling is greatly minimized. The impact is reduced to the degree that the standard deviation of response time is less than 10% of regular Java.

For more information about configuring garbage collection, see *Tuning the IBM virtual machine for Java*.

JVM performance

WebSphere eXtreme Scale can run on different versions of Java Platform, Standard Edition. WebSphere eXtreme Scale supports Java SE Version 1.4.2 and above. For improved developer productivity and performance, use Java SE 5 or later to take advantage of annotations and improved garbage collection. WebSphere eXtreme Scale works on 32 bit or 64 bit Java virtual machines.

WebSphere eXtreme Scale is tested with a subset of the available virtual machines, however, the supported list is not exclusive. You can run WebSphere eXtreme Scale on any vendor JVM at Version 1.4.2 or later. However, if a problem occurs with a vendor JVM, you must contact the JVM vendor for support. If possible, use the JVM from the WebSphere run time on any platform that WebSphere Application Server supports.

For most of the scenarios in which WebSphere eXtreme Scale is used, Java SE Version 6 of the JVM performs better than Edition 5 or 1.4. Java Platform, Standard Edition Version 1.4 performs poorly especially for scenarios that use the **gencon** collector. In general, use the latest available version of Java Platform, Standard Edition for the best performance.

Heap size

The recommendation is 1 to 2 GB heaps with a JVM per four cores. The optimum heap size number depends on the following factors:

- Number of live objects in the heap.
- Complexity of live objects in the heap.
- Number of available cores for the JVM.

For example, an application that stores 10 K byte arrays can run a much larger heap than an application that uses complex graphs of POJOs.

Thread count

The thread count depends on a few factors. A limit exists for how many threads a single shard can manage. A shard is an instance of a partition, and can be a primary or a replica. With more shards for each JVM, you have more threads with each additional shard providing more concurrent paths to the data. Each shard is as concurrent as possible although there is a limit to the concurrency.

Object Request Broker (ORB) requirements

The IBM SDK includes an IBM ORB implementation that has been tested with WebSphere Application Server and WebSphere eXtreme Scale. To ease the support process, use an IBM-provided JVM. Other JVM implementations use a different ORB. The IBM ORB is only supplied with IBM-provided Java virtual machines. WebSphere eXtreme Scale requires a working ORB to operate. You can use WebSphere eXtreme Scale with ORBs from other vendors. However, if you have a problem with a vendor ORB, you must contact the ORB vendor for support. The

IBM ORB implementation is compatible with third party Java virtual machines and can be substituted if needed.

orb.properties tuning

In the lab, the following file was used on data grids of up to 1500 JVMs. The orb.properties file is in the lib folder of the runtime environment.

```
# IBM JDK properties for ORB
org.omg.CORBA.ORBClass=com.ibm.CORBA.iiop.ORB
org.omg.CORBA.ORBSingletonClass=com.ibm.rmi.corba.ORBSingleton

# WS Interceptors
org.omg.PortableInterceptor.ORBInitializerClass=com.ibm.ws.objectgrid.corba.ObjectGridInitializer

# WS ORB & Plugins properties
com.ibm.CORBA.ForceTunnel=never
com.ibm.CORBA.RequestTimeout=10
com.ibm.CORBA.ConnectTimeout=10

# Needed when lots of JVMs connect to the catalog at the same time
com.ibm.CORBA.ServerSocketQueueDepth=2048

# Clients and the catalog server can have sockets open to all JVMs
com.ibm.CORBA.MaxOpenConnections=1016

# Thread Pool for handling incoming requests, 200 threads here
com.ibm.CORBA.ThreadPool.IsGrowable=false
com.ibm.CORBA.ThreadPool.MaximumSize=200
com.ibm.CORBA.ThreadPool.MinimumSize=200
com.ibm.CORBA.ThreadPool.InactivityTimeout=180000

# No splitting up large requests/responses in to smaller chunks
com.ibm.CORBA.FragmentSize=0
```

CopyMode best practices

WebSphere eXtreme Scale makes a copy of the value based on the six available CopyMode settings. Determine which setting works best for your deployment requirements.

You can use the BackingMap API `setCopyMode(CopyMode, valueInterfaceClass)` method to set the copy mode to one of the following final static fields that are defined in the `com.ibm.websphere.objectgrid.CopyMode` class.

When an application uses the ObjectMap interface to obtain a reference to a map entry, use that reference only within the WebSphere eXtreme Scale transaction that obtained the reference. Using the reference in a different transaction can lead to errors. For example, if you use the pessimistic locking strategy for the BackingMap, a `get` or `getForUpdate` method call acquires an S (shared) or U (update) lock, depending on the transaction. The `get` method returns the reference to the value and the lock that is obtained is released when the transaction completes. The transaction must call the `get` or `getForUpdate` method to lock the map entry in a different transaction. Each transaction must obtain its own reference to the value by calling the `get` or `getForUpdate` method instead of reusing the same value reference in multiple transactions.

CopyMode for entity maps

When using a map associated with an EntityManager API entity, the map always returns the entity Tuple objects directly without making a copy unless you are using COPY_TO_BYTES copy mode. It is important that the CopyMode is updated or the Tuple is copied appropriately when making changes.

COPY_ON_READ_AND_COMMIT

The `COPY_ON_READ_AND_COMMIT` mode is the default mode. The `valueInterfaceClass` argument is ignored when this mode is used. This mode ensures that an application does not contain a reference to the value object that is in the `BackingMap`. Instead, the application is always working with a copy of the value that is in the `BackingMap`. The `COPY_ON_READ_AND_COMMIT` mode ensures that the application can never inadvertently corrupt the data that is cached in the `BackingMap`. When an application transaction calls an `ObjectMap.get` method for a given key, and it is the first access of the `ObjectMap` entry for that key, a copy of the value is returned. When the transaction is committed, any changes that are committed by the application are copied to the `BackingMap` to ensure that the application does not have a reference to the committed value in the `BackingMap`.

COPY_ON_READ

The `COPY_ON_READ` mode improves performance over the `COPY_ON_READ_AND_COMMIT` mode by eliminating the copy that occurs when a transaction is committed. The `valueInterfaceClass` argument is ignored when this mode is used. To preserve the integrity of the `BackingMap` data, the application ensures that every reference that it has for an entry is destroyed after the transaction is committed. With this mode, the `ObjectMap.get` method returns a copy of the value instead of a reference to the value to ensure that changes that are made by the application to the value does not affect the `BackingMap` value until the transaction is committed. However, when the transaction does commit, a copy of changes is not made. Instead, the reference to the copy that was returned by the `ObjectMap.get` method is stored in the `BackingMap`. The application destroys all map entry references after the transaction is committed. If application does not destroy the map entry references, the application might cause the data cached in `BackingMap` to become corrupted. If an application is using this mode and is having problems, switch to `COPY_ON_READ_AND_COMMIT` mode to see if the problem still exists. If the problem goes away, then the application is failing to destroy all of its references after the transaction has committed.

COPY_ON_WRITE

The `COPY_ON_WRITE` mode improves performance over the `COPY_ON_READ_AND_COMMIT` mode by eliminating the copy that occurs when the `ObjectMap.get` method is called for the first time by a transaction for a given key. The `ObjectMap.get` method returns a proxy to the value instead of a direct reference to the value object. The proxy ensures that a copy of the value is not made unless the application calls a set method on the value interface that is specified by the `valueInterfaceClass` argument. The proxy provides a copy on write implementation. When a transaction commits, the `BackingMap` examines the proxy to determine if any copy was made as a result of a set method being called. If a copy was made, then the reference to that copy is stored in the `BackingMap`. The big advantage of this mode is that a value is never copied on a read or at a commit when the transaction never calls a set method to change the value.

The `COPY_ON_READ_AND_COMMIT` and `COPY_ON_READ` modes both make a deep copy when a value is retrieved from the `ObjectMap`. If an application only updates some of the values that are retrieved in a transaction then this mode is not optimal. The `COPY_ON_WRITE` mode supports this behavior in an efficient manner but requires that the application uses a simple pattern. The value objects are required to support an interface. The application must use the methods on this

interface when it is interacting with the value in an eXtreme Scale Session. If this is the case, then eXtreme Scale creates proxies for the values that are returned to the application. The proxy has a reference to the real value. If the application performs read operations only, the read operations always run against the real copy. If the application modifies an attribute on the object, the proxy makes a copy of the real object and then makes the modification on the copy. The proxy then uses the copy from that point on. Using the copy allows the copy operation to be avoided completely for objects that are only read by the application. All modify operations must start with the set prefix. Enterprise JavaBeans normally are coded to use this style of method naming for methods that modify the objects attributes. This convention must be followed. Any objects that are modified are copied at the time that they are modified by the application. This read and write scenario is the most efficient scenario supported by eXtreme Scale. To configure a map to use COPY_ON_WRITE mode, use the following example. In this example, the application stores Person objects that are keyed using the name in the Map. The person object is represented in the following code snippet.

```
class Person {
    String name;
    int age;
    public Person() {
    }
    public void setName(String n) {
        name = n;
    }
    public String getName() {
        return name;
    }
    public void setAge(int a) {
        age = a;
    }
    public int getAge() {
        return age;
    }
}
```

The application uses the IPerson interface only when it interacts with values that are retrieved from a ObjectMap. Modify the object to use an interface as in the following example.

```
interface IPerson
{
    void setName(String n);
    String getName();
    void setAge(int a);
    int getAge();
}
// Modify Person to implement IPerson interface
class Person implements IPerson {
    ...
}
```

The application then needs to configure the BackingMap to use COPY_ON_WRITE mode, like in the following example:

```
ObjectGrid dg = ...;
BackingMap bm = dg.defineMap("PERSON");
// use COPY_ON_WRITE for this Map with
// IPerson as the valueProxyInfo Class
bm.setCopyMode(CopyMode.COPY_ON_WRITE,IPerson.class);
// The application should then use the following
// pattern when using the PERSON Map.
Session sess = ...;
ObjectMap person = sess.getMap("PERSON");
```

```

...
sess.begin();
// the application casts the returned value to IPerson and not Person
IPerson p = (IPerson)person.get("Billy");
p.setAge(p.getAge()+1);
...
// make a new Person and add to Map
Person p1 = new Person();
p1.setName("Bobby");
p1.setAge(12);
person.insert(p1.getName(), p1);
sess.commit();
// the following snippet WON'T WORK. Will result in ClassCastException
sess.begin();
// the mistake here is that Person is used rather than
// IPerson
Person a = (Person)person.get("Bobby");
sess.commit();

```

The first section shows the application retrieving a value that was named Billy in the map. The application casts the returned value to the IPerson object, not the Person object because the proxy that is returned implements two interfaces:

- The interface specified in the BackingMap.setCopyMode method call
- The com.ibm.websphere.objectgrid.ValueProxyInfo interface

You can cast the proxy to two types. The last part of the preceding code snippet demonstrates what is not allowed in COPY_ON_WRITE mode. The application retrieves the Bobby record and tries to cast the record to a Person object. This action fails with a class cast exception because the proxy that is returned is not a Person object. The returned proxy implements the IPerson object and ValueProxyInfo.

ValueProxyInfo interface and partial update support: This interface allows an application to retrieve either the committed read-only value referenced by the proxy or the set of attributes that have been modified during this transaction.

```

public interface ValueProxyInfo {
    List /**/ ibmGetDirtyAttributes();
    Object ibmGetRealValue();
}

```

The ibmGetRealValue method returns a read only copy of the object. The application must not modify this value. The ibmGetDirtyAttributes method returns a list of strings representing the attributes that have been modified by the application during this transaction. The main use case for ibmGetDirtyAttributes is in a Java database connectivity (JDBC) or CMP based loader. Only the attributes that are named in the list need be updated on either the SQL statement or object mapped to the table, which leads to more efficient SQL generated by the Loader. When a copy on write transaction is committed and if a loader is plugged in, the loader can cast the values of the modified objects to the ValueProxyInfo interface to obtain this information.

Handling the equals method when using COPY_ON_WRITE or proxies: For example, the following code constructs a Person object and then inserts it to an ObjectMap. Next, it retrieves the same object using the ObjectMap.get method. The value is cast to the interface. If the value is cast to the Person interface, a ClassCastException exception results because the returned value is a proxy that implements the IPerson interface and is not a Person object. The equality check fails when using the == operation because they are not the same object.

```

session.begin();
// new the Person object
Person p = new Person(...);
personMap.insert(p.getName, p);
// retrieve it again, remember to use the interface for the cast
IPerson p2 = personMap.get(p.getName());
if(p2 == p) {
    // they are the same
} else {
    // they are not
}

```

Another consideration is when you must override the equals method. As illustrated in the following snippet of code, the equals method must verify that the argument is an object that implements IPerson interface and cast the argument to be a IPerson. Because the argument might be a proxy that implements the IPerson interface, you must use the getAge and getName methods when comparing instance variables for equality.

```

{
    if ( obj == null ) return false;
    if ( obj instanceof IPerson ) {
        IPerson x = (IPerson) obj;
        return ( age.equals( x.getAge() ) && name.equals( x.getName() ) )
    }
    return false;
}

```

ObjectQuery and HashIndex configuration requirements: When using COPY_ON_WRITE with ObjectQuery or a HashIndex plug-in, it's important to configure the ObjectQuery schema and HashIndex plug-in to access the objects using property methods, which is the default. If configured to use field access, the query engine and index will attempt to access the fields in the proxy object, which will always return null or 0 since the object instance will be a proxy.

NO_COPY

The NO_COPY mode allows an application to ensure that it never modifies a value object that is obtained using an ObjectMap.get method in exchange for performance improvements. The valueInterfaceClass argument is ignored when this mode is used. If this mode is used, no copy of the value is ever made. If the application modifies values, then the data in the BackingMap is corrupted. The NO_COPY mode is primarily useful for read-only maps where data is never modified by the application. If the application is using this mode and it is having problems, then switch to the COPY_ON_READ_AND_COMMIT mode to see if the problem still exists. If the problem goes away, then the application is modifying the value returned by ObjectMap.get method, either during transaction or after transaction has committed. All maps associated with EntityManager API entities automatically use this mode regardless of what is specified in the eXtreme Scale configuration.

All maps associated with EntityManager API entities automatically use this mode regardless of what is specified in the eXtreme Scale configuration.

COPY_TO_BYTES

You can store objects in a serialized format instead of POJO format. By using the COPY_TO_BYTES setting, you can reduce the memory footprint that a large graph of Objects can consume. See “Byte array maps” on page 41 for additional information.

Incorrect use of CopyMode

Errors occur when an application attempts to improve performance by using the `COPY_ON_READ`, `COPY_ON_WRITE`, or `NO_COPY` copy mode, as described above. The intermittent errors do not occur when you change the copy mode to the `COPY_ON_READ_AND_COMMIT` mode.

Problem

The problem might be due to corrupted data in the ObjectGrid map, which is a result of the application violating the programming contract of the copy mode that is being used. Data corruption can cause unpredictable errors to occur intermittently or in an unexplained or unexpected fashion.

Solution

The application must comply with the programming contract that is stated for the copy mode being used. For the `COPY_ON_READ` and `COPY_ON_WRITE` copy modes, the application uses a reference to a value object outside of the transaction scope from which the value reference was obtained. To use these modes, the application must delete the reference to the value object after the transaction completes, and obtain a new reference to the value object in each transaction that accesses the value object. For the `NO_COPY` copy mode, the application must never change the value object. In this case, either write the application so that it does not change the value object, or set the application to use a different copy mode.

Byte array maps

You can store the key-value pairs in your maps in a byte array instead of POJO form, which reduces the memory footprint that a large graph of objects can consume.

Advantages

The amount of memory that is consumed increases with the number of objects in a graph of objects. By reducing a complicated graph of objects to a byte array, only one object is maintained in the heap instead of several objects. With this reduction of the number of objects in the heap, the Java run time has fewer objects to search for during garbage collection.

The default copy mechanism used by WebSphere eXtreme Scale is serialization, which is expensive. For instance, if using the default copy mode of `COPY_ON_READ_AND_COMMIT`, a copy is made both at read time and at get time. Instead of making a copy at read time, with byte arrays, the value is inflated from bytes, and instead of making a copy at commit time, the value is serialized to bytes. Using byte arrays results in equivalent data consistency to the default setting with a reduction of memory used.

When using byte arrays, note that having an optimized serialization mechanism is critical to seeing a reduction of memory consumption. For more information, see “Serialization performance” on page 214.

Configuring byte array maps

You can enable byte array maps with the ObjectGrid XML file by modifying the CopyMode attribute that is used by a map to the setting COPY_TO_BYTES, shown in the following example:

```
<backingMap name="byteMap" copyMode="COPY_TO_BYTES" />
```

See the topic on the ObjectGrid descriptor XML file in the *Administration Guide* for more information.

Considerations

You must consider whether or not to use byte array maps in a given scenario. Although you can reduce your memory use, processor use can increase when you use byte arrays.

The following list outlines several factors that should be considered before choosing to use the byte array map function.

Object type

Comparatively, memory reduction may not be possible when using byte array maps for some object types. Consequently, several types of objects exist for which you should not use byte array maps. If you are using any of the Java primitive wrappers as values, or a POJO that does not contain references to other objects (only storing primitive fields), the number of Java Objects is already as low as possible—there is only one. Since the amount of memory used by the object is already optimized, using a byte array map for these types of objects is not recommended. Byte array maps are more suitable to object types that contain other objects or collections of objects where the total number of POJO objects is greater than one.

For example, if you have a Customer object that had a business Address and a home Address, as well as a collection of Orders, the number of objects in the heap and the number of bytes used by those objects can be reduced by using byte array maps.

Local access

When using other copy modes, applications can be optimized when copies are made if objects are Cloneable with the default ObjectTransformer or when a custom ObjectTransformer is provided with an optimized copyValue method. Compared to the other copy modes, copying on reads, writes, or commit operations will have additional cost when accessing objects locally. For example, if you have a near cache in a distributed topology or are directly accessing a local or server ObjectGrid instance, the access and commit time will increase when using byte array maps due to the cost of serialization. You will see a similar cost in a distributed topology if you use data grid agents or you access the server primary when using the ObjectGridEventGroup.ShardEvents plug-in.

Plug-in interactions

With byte array maps, objects are not inflated when communicating from a client to a server unless the server needs the POJO form. Plug-ins that interact with the map value will experience a reduction in performance due to the requirement to inflate the value.

Any plug-in that uses `LogElement.getCacheEntry` or `LogElement.getCurrentValue` will see this additional cost. If you want to get the key, you can use `LogElement.getKey`, which avoids the additional overhead associated with the `LogElement.getCacheEntry().getKey` method. The following sections discuss plug-ins in light of the usage of byte arrays.

Indexes and queries

When objects are stored in POJO format, the cost of doing indexing and querying is minimal because the object does not need to be inflated. When using a byte array map you will have the additional cost of inflating the object. In general if your application uses indexes or queries, it is not recommended to use byte array maps unless you only run queries on key attributes.

Optimistic locking

When using the optimistic locking strategy, you will have the additional cost during updates and invalidate operations. This comes from having to inflate the value on the server to get the version value to do optimistic collision checking. If you are just using optimistic locking to guarantee fetch operations and do not need optimistic collision checking, you can use the `com.ibm.websphere.objectgrid.plugins.builtins.NoVersioningOptimisticCallback` to disable version checking.

Loader

With a Loader, you will also have the cost in the eXtreme Scale run time from inflating and reserializing the value when it is used by the Loader. You can still use byte array maps with Loaders, but consider the cost of making changes to the value in such a scenario. For example, you can use the byte array feature in the context of a read mostly cache. In this case, the benefit of having less objects in the heap and less memory used will outweigh the cost incurred from using byte arrays on insert and update operations.

ObjectGridEventListener

When using the `transactionEnd` method in the `ObjectGridEventListener` plug-in, you will have an additional cost on the server side for remote requests when accessing a `LogElement's` `CacheEntry` or current value. If the implementation of the method does not access these fields, then you will not have the additional cost.

Plug-in evictor performance best practices

If you use plug-in evictors, they are not active until you create them and associate them with a backing map. The following best practices will increase performance for least frequently used (LFU) and least recently used (LRU) evictors.

Least frequently used (LFU) evictor

The concept of a LFU evictor is to remove entries from the map that are used infrequently. The entries of the map are spread over a set amount of binary heaps. As the usage of a particular cache entry grows, it becomes ordered higher in the heap. When the evictor attempts a set of evictions it removes only the cache entries that are located lower than a specific point on the binary heap. As a result, the least frequently used entries are evicted.

Least recently used (LRU) evictor

The LRU Evictor follows the same concepts of the LFU Evictor with a few differences. The main difference is that the LRU uses a first in, first out queue (FIFO) instead of a set of binary heaps. Every time a cache entry is accessed, it moves to the head of the queue. Consequently, the front of the queue contains the most recently used map entries and the end becomes the least recently used map entries. For example, the A cache entry is used 50 times, and the B cache entry is used only once right after the A cache entry. In this situation, the B cache entry is at the front of the queue because it was used most recently, and the A cache entry is at the end of the queue. The LRU evictor evicts the cache entries that are at the tail of the queue, which are the least recently used map entries.

LFU and LRU properties and best practices to improve performance

Number of heaps

When using the LFU evictor, all of the cache entries for a particular map are ordered over the number of heaps that you specify, improving performance drastically and preventing all of the evictions from synchronizing on one binary heap that contains all of the ordering for the map. More heaps also speeds up the time that is required for reordering the heaps because each heap has fewer entries. Set the number of heaps to 10% of the number of entries in your BaseMap.

Number of queues

When using the LRU evictor, all of the cache entries for a particular map are ordered over the number of LRU queues that you specify, improving performance drastically and preventing all of the evictions from synchronizing on one queue that contains all of the ordering for the map. Set the number of queues to 10% of the number of entries in your BaseMap.

MaxSize property

When an LFU or LRU evictor begins evicting entries, it uses the MaxSize evictor property to determine how many binary heaps or LRU queue elements to evict. For example, assume that you set the number of heaps or queues to have about 10 map entries in each map queue. If your MaxSize property is set to 7, the evictor evicts 3 entries from each heap or queue object to bring the size of each heap or queue back down to 7. The evictor only evicts map entries from a heap or queue when that heap or queue has more than the MaxSize property value of elements in it. Set the MaxSize to 70% of your heap or queue size. For this example, the value is set to 7. You can get an approximate size of each heap or queue by dividing the number of BaseMap entries by the number of heaps or queues that are used.

SleepTime property

An evictor does not constantly remove entries from your map. Instead it is idle for a set amount of time, only checking the map every n number of seconds, where n refers to the SleepTime property. This property also positively affects performance: running an eviction sweep too often lowers performance because of the resources that are needed for processing them. However, not using the evictor often can result in a map that has entries that are not needed. A map full of entries that are not needed can negatively affect both the memory requirements and processing resources that are required for your map. Setting the eviction sweep interval to

fifteen seconds is a good practice for most maps. If the map is written to frequently and is used at a high transaction rate, consider setting the value to a lower time. If the map is accessed infrequently, you can set the time to a higher value.

Example

The following example defines a map, creates a new LFU evictor, sets the evictor properties, and sets the map to use the evictor:

```
//Use ObjectGridManager to create/get the ObjectGrid. Refer to
// the ObjectGridManger section
ObjectGrid objGrid = ObjectGridManager.create.....
BackingMap bMap = objGrid.defineMap("SomeMap");

//Set properties assuming 50,000 map entries
LFUEvictor someEvictor = new LFUEvictor();
someEvictor.setNumberOfHeaps(5000);
someEvictor.setMaxSize(7);
someEvictor.setSleepTime(15);
bMap.setEvictor(someEvictor);
```

Using the LRU evictor is very similar to using an LFU evictor. An example follows:

```
ObjectGrid objGrid = new ObjectGrid;
BackingMap bMap = objGrid.defineMap("SomeMap");

//Set properties assuming 50,000 map entries
LRUEvictor someEvictor = new LRUEvictor();
someEvictor.setNumberOfLRUQueues(5000);
someEvictor.setMaxSize(7);
someEvictor.setSleepTime(15);
bMap.setEvictor(someEvictor);
```

Notice that only two lines are different from the LFUEvictor example.

Locking performance best practices

Locking strategies and transaction isolation settings affect the performance of your applications.

Retrieve a cached instance

Pessimistic locking strategy

Use the pessimistic locking strategy for read and write map operations where keys often collide. The pessimistic locking strategy has the greatest impact on performance.

Read committed and read uncommitted transaction isolation

When you are using pessimistic locking strategy, set the transaction isolation level using the `Session.setTransactionIsolation` method. For read committed or read uncommitted isolation, use the `Session.TRANSACTION_READ_COMMITTED` or `Session.TRANSACTION_READ_UNCOMMITTED` arguments depending on the isolation. To reset the transaction isolation level to the default pessimistic locking behavior, use the `Session.setTransactionIsolation` method with the `Session.REPEATABLE_READ` argument.

Read committed isolation reduces the duration of shared locks, which can improve concurrency and reduce the chance for deadlocks. This isolation level should be

used when a transaction does not need assurances that read values remain unchanged for the duration of the transaction.

Use an uncommitted read when the transaction does not need to see the committed data.

Optimistic locking strategy

Optimistic locking is the default configuration. This strategy improves both performance and scalability compared to the pessimistic strategy. Use this strategy when your applications can tolerate some optimistic update failures, while still performing better than the pessimistic strategy. This strategy is excellent for read operations and infrequent update applications.

OptimisticCallback plug-in

The optimistic locking strategy makes a copy of the cache entries and compares them as needed. This operation can be expensive because copying the entry might involve cloning or serialization. To implement the fastest possible performance, implement the custom plug-in for non-entity maps.

See for more information. See the information about the OptimisticCallback plug-in in the *Product Overview* for more information.

Use version fields for entities

When you are using optimistic locking with entities, use the @Version annotation or the equivalent attribute in the Entity metadata descriptor file. The version annotation gives the ObjectGrid a very efficient way of tracking the version of an object. If the entity does not have a version field and optimistic locking is used for the entity, then the entire entity must be copied and compared.

None locking strategy

Use the none locking strategy for applications that are read only. The none locking strategy does not obtain any locks or use a lock manager. Therefore, this strategy offers the most concurrency, performance and scalability.

Serialization performance

WebSphere eXtreme Scale uses multiple Java processes to hold data. These processes serialize the data: That is, they convert the data (which is in the form of Java object instances) to bytes and back to objects again as needed to move the data between client and server processes. Marshalling the data is the most expensive operation and must be addressed by the application developer when designing the schema, configuring the data grid and interacting with the data-access APIs.

The default Java serialization and copy routines are relatively slow and can consume 60 to 70 percent of the processor in a typical setup. The following sections are choices for improving the performance of the serialization.

Write an ObjectTransformer for each BackingMap

An ObjectTransformer can be associated with a BackingMap. Your application can have a class that implements the ObjectTransformer interface and provides implementations for the following operations:

- Copying values
- Serializing and inflating keys to and from streams
- Serializing and inflating values to and from streams

The application does not need to copy keys because keys are considered immutable.

For more information, see Plug-ins for serializing and copying cached objects and ObjectTransformer interface best practices.

Note: The ObjectTransformer is only invoked when the ObjectGrid knows about the data that is being transformed. For example, when DataGrid API agents are used, the agents themselves as well as the agent instance data or data returned from the agent must be optimized using custom serialization techniques. The ObjectTransformer is not invoked for DataGrid API agents.

Using entities

When using the EntityManager API with entities, the ObjectGrid does not store the entity objects directly into the BackingMaps. The EntityManager API converts the entity object to Tuple objects. See For more information, see the topic on using a loader with entity maps and tuples in the *Programming Guide*. Entity maps are automatically associated with a highly optimized ObjectTransformer. Whenever the ObjectMap API or EntityManager API is used to interact with entity maps, the entity ObjectTransformer is invoked.

Custom serialization

There are some cases when objects must be modified to use custom serialization, such as implementing the `java.io.Externalizable` interface or by implementing the `writeObject` and `readObject` methods for classes implementing the `java.io.Serializable` interface. Custom serialization techniques should be employed when the objects are serialized using mechanisms other than the ObjectGrid API or EntityManager API methods.

For example, when objects or entities are stored as instance data in a DataGrid API agent or the agent returns objects or entities, those objects are not transformed using an ObjectTransformer. The agent, will however, automatically use the ObjectTransformer when using `EntityMixin` interface. See DataGrid agents and entity based Maps for further details.

Byte arrays

When using the ObjectMap or DataGrid APIs, the key and value objects are serialized whenever the client interacts with the data grid and when the objects are replicated. To avoid the overhead of serialization, use byte arrays instead of Java objects. Byte arrays are much cheaper to store in memory since the JDK has less objects to search for during garbage collection and they are can be inflated only

when needed. Byte arrays should only be used if you do not need to access the objects using queries or indexes. Since the data is stored as bytes, the data can only be accessed through its key.

WebSphere eXtreme Scale can automatically store data as byte arrays using the `CopyMode.COPY_TO_BYTES` map configuration option, or it can be handled manually by the client. This option will store the data efficiently in memory and can also automatically inflate the objects within the byte array for use by query and indexes on demand.

ObjectTransformer interface best practices

The `ObjectTransformer` interface uses callbacks to the application to provide custom implementations of common and expensive operations such as object serialization and deep copies on objects.

Overview

For details about the `ObjectTransformer` interface, see “`ObjectTransformer` plug-in” on page 210. From a performance viewpoint, and from the `CopyMode` method information that is in the `CopyMode` method best practices topic, eXtreme Scale clearly copies the values for all cases except when `NO_COPY` mode is used. The default copying mechanism that is employed in eXtreme Scale is serialization, which is known as an expensive operation. The `ObjectTransformer` interface is used in this situation. The `ObjectTransformer` interface uses callbacks to the application to provide a custom implementation of common and expensive operations, such as object serialization and deep copies on objects.

An application can provide an implementation of the `ObjectTransformer` interface to a map, and eXtreme Scale then delegates to the methods on this object and relies on the application to provide an optimized version of each method in the interface. The `ObjectTransformer` interface follows:

```
public interface ObjectTransformer {
    void serializeKey(Object key, ObjectOutputStream stream) throws IOException;
    void serializeValue(Object value, ObjectOutputStream stream) throws IOException;
    Object inflateKey(ObjectInputStream stream) throws IOException, ClassNotFoundException;
    Object inflateValue(ObjectInputStream stream) throws IOException, ClassNotFoundException;
    Object copyValue(Object value);
    Object copyKey(Object key);
}
```

You can associate an `ObjectTransformer` interface with a `BackingMap` by using the following example code:

```
ObjectGrid g = ...;
BackingMap bm = g.defineMap("PERSON");
MyObjectTransformer ot = new MyObjectTransformer();
bm.setObjectTransformer(ot);
```

Tune object serialization and inflation

Object serialization is typically the most important performance consideration with eXtreme Scale, which uses the default serializable mechanism if an `ObjectTransformer` plug-in is not supplied by the application. An application can provide implementations of either the `Serializable` `readObject` and `writeObject`, or it can have the objects implement the `Externalizable` interface, which is approximately ten times faster. If the objects in the map cannot be modified, then an application can associate an `ObjectTransformer` interface with the `ObjectMap`. The `serialize` and `inflate` methods are provided to allow the application to provide custom code to optimize these operations, given their large performance impact on

the system. The `serialize` method serializes the object to the provided stream. The `inflate` method provides the input stream and expects the application to create the object, inflate it using data in the stream and return the object. Implementations of the `serialize` and `inflate` methods must mirror each other.

Tune deep copy operations

After an application receives an object from an `ObjectMap`, eXtreme Scale performs a deep copy on the object value to ensure that the copy in the `BaseMap` map maintains data integrity. The application can then modify the object value safely. When the transaction commits, the copy of the object value in the `BaseMap` map is updated to the new modified value and the application stops using the value from that point on. You could have copied the object again at the commit phase to make a private copy. However, in this case the performance cost of this action was traded off against requiring the application programmer not to use the value after the transaction commits. The default `ObjectTransformer` attempts to use either a clone or a `serialize` and `inflate` pair to generate a copy. The `serialize` and `inflate` pair is the worst case performance scenario. If profiling reveals that `serialize` and `inflate` is a problem for your application, write an appropriate clone method to create a deep copy. If you cannot alter the class, then create a custom `ObjectTransformer` plug-in and implement more efficient `copyValue` and `copyKey` methods.

Query performance tuning

To tune the performance of your queries, use the following techniques and tips.

Using parameters

When a query runs, the query string must be parsed and a plan developed to run the query, both of which can be costly. WebSphere eXtreme Scale caches query plans by the query string. Since the cache is a finite size, it is important to reuse query strings whenever possible. Using named or positional parameters also helps performance by fostering query plan reuse.

```
Positional Parameter Example Query q = em.createQuery("select c from Customer c where c.surname=?1"); q.setParameter(1, "Claus");
```

Using indexes

Proper indexing on a map might have a significant impact on query performance, even though indexing has some overhead on overall map performance. Without indexing on object attributes involved in queries, the query engine performs a table scan for each attribute. The table scan is the most expensive operation during a query run. Indexing on object attributes that are involved in queries allow the query engine to avoid an unnecessary table scan, improving the overall query performance. If the application is designed to use query intensively on a read-most map, configure indexes for object attributes that are involved in the query. If the map is mostly updated, then you must balance between query performance improvement and indexing overhead on the map. See [Indexing](#) for more information.

When plain old Java objects (POJO) are stored in a map, proper indexing can avoid a Java reflection. In the following example, query replaces the `WHERE` clause with range index search, if the `budget` field has an index built over it. Otherwise, query scans the entire map and evaluates the `WHERE` clause by first getting the budget using Java reflection and then comparing the budget with the value 50000:

```
SELECT d FROM DeptBean d WHERE d.budget=50000
```

See “Query plan” on page 117 for details on how to best tune individual queries and how different syntax, object models and indexes can affect query performance.

Using pagination

In client-server environments, the query engine transports the entire result map to the client. The data that is returned should be divided into reasonable chunks. The EntityManager Query and ObjectMap ObjectQuery interfaces both support the setFirstResult and setMaxResults methods that allow the query to return a subset of the results.

Return primitive values instead of entities

With the EntityManager Query API, entities are returned as query parameters. The query engine currently returns the keys for these entities to the client. When the client iterates over these entities using the Iterator from the getResultIterator method, each entity is automatically inflated and managed as if it were created with the find method on the EntityManager interface. The entire entity graph is built from the entity ObjectMap on the client. The entity value attributes and any related entities are eagerly resolved.

To avoid building the costly graph, modify the query to return the individual attributes with path navigation.

For example:

```
// Returns an entity  
SELECT p FROM Person p  
// Returns attributes SELECT p.name, p.address.street, p.address.city, p.gender FROM Person p
```

Chapter 10. Troubleshooting

In addition to the logs and trace, messages, and release notes discussed in this section, you can use monitoring tools to figure out issues such as the location of data in the environment, the availability of servers in the data grid, and so on. If you are running in a WebSphere Application Server environment, you can use Performance Monitoring Infrastructure (PMI). If you are running in a stand-alone environment, you can use a vendor monitoring tool, such as CA Wily Introscope or Hyperic HQ. You can also use and customize the xsAdmin sample utility to display textual information about your environment.

Enabling logging

You can use logs to monitor and troubleshoot your environment.

About this task

Logs are saved different locations and formats depending on your configuration.

Procedure

- **Enable logs in a stand-alone environment.**

With stand-alone catalog servers, the logs are in the location where you run the **startOgServer** command. For container servers, you can use the default location or set a custom log location:

- **Default log location:** The logs are in the directory where the server command was run. If you start the servers in the `wxs_home/bin` directory, the logs and trace files are in the `logs/<server_name>` directories in the `bin` directory.
- **Custom log location:** To specify an alternate location for container server logs, create a properties file, such as `server.properties`, with the following contents:

```
workingDirectory=<directory>
traceSpec=
systemStreamToFileEnabled=true
```

The **workingDirectory** property is the root directory for the logs and optional trace file. WebSphere eXtreme Scale creates a directory with the name of the container server with a `SystemOut.log` file, a `SystemErr.log` file, and a trace file. To use a properties file during container startup, use the **-serverProps** option and provide the server properties file location.

- **Enable logs in WebSphere Application Server.**

See WebSphere Application Server: Enabling and disabling logging for more information.

- **Retrieve FFDC files.**

FFDC files are for IBM support to aid in debug. These files might be requested by IBM support if a problem occurs. These files are in a directory labeled, `ffdc`, and contain files that resemble the following:

```
server2_exception.log
server2_20802080_07.03.05_10.52.18_0.txt
```

What to do next

View the log files in their specified locations. Common messages to look for in the `SystemOut.log` file are start confirmation messages, such as the following example:

CWOBJ1001I: ObjectGrid Server catalogServer01 is ready to process requests.

For more information about a specific message in the log files, see “Messages” on page 369.

Collecting trace

You can use trace to monitor and troubleshoot your environment. You must provide trace for a server when you work with IBM support.

About this task

Collecting trace can help you monitor and fix problems in your deployment of WebSphere eXtreme Scale. How you collect trace depends on your configuration. See “Trace options” on page 359 for a list of the different trace specifications you can collect.

Procedure

- **Collect trace within a WebSphere Application Server environment.**

If your catalog and container servers are in a WebSphere Application Server environment, see WebSphere Application Server: Working with trace for more information.

- **Collect trace with the stand-alone catalog or container server start command.**

You can set trace on a catalog service or container server by using the **-traceSpec** and **-traceFile** parameters with the **startOgServer** command. For example:

```
startOgServer.sh catalogServer -traceSpec ObjectGridPlacement=all-enabled -traceFile /home/user1/logs/trace.log
```

The **-traceFile** parameter is optional. If you do not set a **-traceFile** location, the trace file goes to the same location as the system out log files. For more information about these parameters, see the information about the **startOgServer** script in the *Administration Guide*.

- **Collect trace on the stand-alone catalog or container server with a properties file.**

To collect trace from a properties file, create a file, such as a `server.properties` file, with the following contents:

```
workingDirectory=<directory>  
traceSpec=<trace_specification>  
systemStreamToFileEnabled=true
```

The **workingDirectory** property is the root directory for the logs and optional trace file. If the **workingDirectory** value is not set, the default working directory is the location used to start the servers, such as `wxs_home/bin`. To use a properties file during server startup, use the **-serverProps** parameter with the **startOgServer** command and provide the server properties file location. For more information about the server properties file and how to use the file, see the information about the server properties file in the *Administration Guide*.

- **Collect trace on a stand-alone client.**

You can start trace collection on a stand-alone client by adding system properties to the startup script for the client application. In the following example, trace settings are specified for the `com.ibm.samples.MyClientProgram` application:

```
java -DtraceSettingsFile=MyTraceSettings.properties  
-Djava.util.logging.manager=com.ibm.ws.bootstrap.WsLogManager  
-Djava.util.logging.configureByServer=true com.ibm.samples.MyClientProgram
```


See WebSphere Application Server: Enabling trace on client and stand-alone applications for more information.

- **Collect trace with the ObjectGridManager interface.**

You can also set trace during run time on an ObjectGridManager interface. Setting trace on an ObjectGridManager interface can be used to get trace on an eXtreme Scale client while it connects to an eXtreme Scale and commits transactions. To set trace on an ObjectGridManager interface, supply a trace specification and a trace log.

```
ObjectGridManager manager = ObjectGridManagerFactory.getObjectGridManager();
...
manager.setTraceEnabled(true);
manager.setTraceFileName("Logs/myClient.log");
manager.setTraceSpecification("ObjectGridReplication=all=enabled");
```

For more information about the ObjectGridManager interface, see “Interacting with an ObjectGrid using ObjectGridManager” on page 17.

- **Collect trace on container servers with the xsadmin utility.**

To collect trace with the **xsadmin** utility, use the **setTraceSpec** option. Use the **xsadmin** utility to collect trace on a stand-alone environment during run time instead of during startup. You can collect trace on all servers and catalog services or you can filter the servers based on the ObjectGrid name, and other properties. For example, to collect ObjectGridReplication trace with access to the catalog service server, run:

```
xsadmin.bat -setTraceSpec "ObjectGridReplication=all=enabled"
```

You can also disable trace by setting the trace specification to ***=all=disabled**. For more information about the **setTraceSpec** option, see the information about the **xsadmin** utility in the *Administration Guide*.

Results

Trace files are written to the specified location.

Trace options

You can enable trace to provide information about your environment to IBM support.

About trace

WebSphere eXtreme Scale trace is divided into several different components. You can specify the level of trace to use. Common levels of trace include: all, debug, entryExit, and event.

An example trace string follows:

```
ObjectGridComponent=level=enabled
```

You can concatenate trace strings. Use the * (asterisk) symbol to specify a wildcard value, such as **ObjectGrid*=all=enabled**. If you need to provide a trace to IBM support, a specific trace string is requested. For example, if a problem with replication occurs, the **ObjectGridReplication=debug=enabled** trace string might be requested.

Trace specification

ObjectGrid

General core cache engine.

- ObjectGridCatalogServer**
General catalog service.
- ObjectGridChannel**
Static deployment topology communications.
- 7.1+ ObjectGridClientInfo**
DB2 client information.
- 7.1+ ObjectGridClientInfoUser**
DB2 user information.
- ObjectgridCORBA**
Dynamic deployment topology communications.
- ObjectGridDataGrid**
The AgentManager API.
- ObjectGridDynaCache**
The WebSphere eXtreme Scale dynamic cache provider.
- ObjectGridEntityManager**
The EntityManager API. Use with the Projector option.
- ObjectGridEvictors**
ObjectGrid built-in evictors.
- ObjectGridJPA**
Java Persistence API (JPA) loaders.
- ObjectGridJPACache**
JPA cache plug-ins.
- ObjectGridLocking**
ObjectGrid cache entry lock manager.
- ObjectGridMBean**
Management beans.
- 7.1+ ObjectGridMonitor**
Historical monitoring infrastructure.
- ObjectGridPlacement**
Catalog server shard placement service.
- ObjectGridQuery**
ObjectGrid query.
- ObjectGridReplication**
Replication service.
- ObjectGridRouting**
Client/server routing details.
- ObjectGridSecurity**
Security trace.
- ObjectGridStats**
ObjectGrid statistics.
- ObjectGridStreamQuery**
The Stream Query API.
- ObjectGridWriteBehind**
ObjectGrid write behind.

Projector

The engine within the EntityManager API.

QueryEngine

The query engine for the Object Query API and EntityManager Query API.

QueryEnginePlan

Query plan diagnostics.

Troubleshooting client connectivity

There are several common problems specific to clients and client connectivity that you can solve as described in the following sections.

Procedure

Problem: If you are using the EntityManager API or byte array maps with the COPY_TO_BYTES copy mode, client data access methods result in various serialization-related exceptions or a NullPointerException exception.

- The following error occurs when you are using the COPY_TO_BYTES copy mode:

```
java.lang.NullPointerException
  at com.ibm.ws.objectgrid.map.BaseMap$BaseMapObjectTransformer2.inflateObject(BaseMap.java:5278)
  at com.ibm.ws.objectgrid.map.BaseMap$BaseMapObjectTransformer.inflateValue(BaseMap.java:5155)
```

- The following error occurs when you are using the EntityManager API:

```
java.lang.NullPointerException
  at com.ibm.ws.objectgrid.em.GraphTraversalHelper.fluffFetchMD(GraphTraversalHelper.java:323)
  at com.ibm.ws.objectgrid.em.GraphTraversalHelper.fluffFetchMD(GraphTraversalHelper.java:343)
  at com.ibm.ws.objectgrid.em.GraphTraversalHelper.getObjectGraph(GraphTraversalHelper.java:102)
  at com.ibm.ws.objectgrid.ServerCoreEventProcessor.getFromMap(ServerCoreEventProcessor.java:709)
  at com.ibm.ws.objectgrid.ServerCoreEventProcessor.processGetRequest(ServerCoreEventProcessor.java:323)
```

Cause: The EntityManager API and COPY_TO_BYTES copy mode use a metadata repository that is embedded in the data grid. When clients connect, the data grid stores the repository identifiers in the client and caches the identifiers for the duration of the client connection. If you restart the data grid, you lose all metadata and the regenerated identifiers do not match the cached identifiers on the client.

Solution: If you are using the EntityManager API or the COPY_TO_BYTES copy mode, disconnect and reconnect all of the clients if the ObjectGrid is stopped and restarted. Disconnecting and reconnecting the clients refreshes the metadata identifier cache. You can disconnect clients by using the ObjectGridManager.disconnect method or the ObjectGrid.destroy method.

Troubleshooting loaders

Use this information to troubleshoot issues with your database loaders.

Procedure

- **Problem:** When you are using an OpenJPA loader with DB2 in WebSphere Application Server, a closed cursor exception occurs.

The following exception is from DB2 in the

org.apache.openjpa.persistence.PersistenceException log file:

```
[jcc][t4][10120][10898][3.57.82] Invalid operation: result set is closed.
```

Solution: By default, the application server configures the resultSetHoldability custom property with a value of 2 (CLOSE_CURSORS_AT_COMMIT). This property causes DB2 to close its resultSet/cursor at transaction boundaries. To remove the exception, change the value of the custom property to 1

(HOLD_CURSORS_OVER_COMMIT). Set the resultSetHoldability custom property on the following path in the WebSphere Application Server cell: **Resources > JDBC provider > DB2 Universal JDBC Driver Provider > DataSources > data_source_name > Custom properties > New.**

- **Problem** DB2 displays an exception: The current transaction has been rolled back because of a deadlock or timeout. Reason code "2".. SQLCODE=-911, SQLSTATE=40001, DRIVER=3.50.152

This exception occurs because of a lock contention problem when you are running with OpenJPA with DB2 in WebSphere Application Server. The default isolation level for WebSphere Application Server is Repeatable Read (RR), which obtains long-lived locks with DB2.**Solution:**

Set the isolation level to Read Committed to reduce the lock contention. Set the webSphereDefaultIsolationLevel data source custom property to set the isolation level to 2(TRANSACTION_READ_COMMITTED) on the following path in the WebSphere Application Server cell: **Resources > JDBC provider > JDBC_provider > Data sources > data_source_name > Custom properties > New.** For more information about the webSphereDefaultIsolationLevel custom property and transaction isolation levels, see Requirements for setting data access isolation levels.

- **Problem:** When you are using the preload function of the JPAloader or JPAEntityLoader, the following CWOBJ1511 message does not display for the partition in a container server: CWOBJ1511I: GRID_NAME:MAPSET_NAME:PARTITION_ID (primary) is open for business.

Instead, a TargetNotAvailableException exception occurs in the container server, which activates the partition that is specified by the preloadPartition property.

Solution: Set the preloadMode attribute to true if you use a JPAloader or JPAEntityLoader to preload data into the map. If the preloadPartition property of the JPAloader and JPAEntityLoader is set to a value between 0 and total_number_of_partitions - 1, then the JPAloader and JPAEntityLoader try to preload the data from backend database into the map. The following snippet of code illustrates how the preloadMode attribute is set to enable asynchronous preload:

```
BackingMap bm = og.defineMap( "map1" );  
bm.setPreloadMode( true );
```

You can also set the preloadMode attribute by using an XML file as illustrated in the following example:

```
<backingMap name="map1" preloadMode="true" pluginCollectionRef="map1"  
lockStrategy="OPTIMISTIC" />
```

Troubleshooting deadlocks

The following sections describe some of the most common deadlock scenarios and suggestions on how to avoid them.

Before you begin

Implement exception handling in your application. See "Implementing exception handling in locking scenarios" on page 144 for more information.

The following exception displays as a result:

```
com.ibm.websphere.objectgrid.plugins.LockDeadlockException: Message
```

This message represents the string that is passed as a parameter when the exception is created and thrown.

Procedure

- **Problem:** LockTimeoutException exception.

Description: When a transaction or client asks for a lock to be granted for a specific map entry, the request often waits for the current client to release the lock before the request is submitted. If the lock request remains idle for an extended period of time, and a lock is never granted, LockTimeoutException exception is created to prevent a deadlock, which is described in more detail in the following section. You are more likely to see this exception when using a pessimistic locking strategy, because the lock never releases until the transaction commits.

Retrieve more details:

The LockTimeoutException exception contains the getLockRequestQueueDetails method, which returns a string. You can use this method to see a detailed description of the situation that triggers the exception. The following is an example of code that catches the exception, and displays an error message.

```
try {
    ...
}
catch (LockTimeoutException lte) {
    System.out.println(lte.getLockRequestQueueDetails());
}
```

The output result is:

```
lock request queue
->[TX:163C269E-0105-4000-E0D7-5B3B090A571D, state =
    Granted 5348 milli-seconds ago, mode = U]
->[TX:163C2734-0105-4000-E024-5B3B090A571D, state =
    Waiting for 5348 milli-seconds, mode = U]
->[TX:163C328C-0105-4000-E114-5B3B090A571D, state =
    Waiting for 1402 milli-seconds, mode = U]
```

If you receive the exception in an ObjectGridException exception catch block, the following code determines the exception and displays the queue details. It also uses the findRootCause utility method.

```
try {
    ...
}
catch (ObjectGridException oe) {
    Throwable Root = findRootCause( oe );
    if (Root instanceof LockTimeoutException) {
        LockTimeoutException lte = (LockTimeoutException)Root;
        System.out.println(lte.getLockRequestQueueDetails());
    }
}
```

Solution: A LockTimeoutException exception prevents possible deadlocks in your application. An exception of this type results when the exception waits a set amount of time. You can set the amount of time that the exception waits by using the setLockTimeout(int) method, which is available for the BackingMap. If a deadlock does not actually exist in your application, adjust the lock timeout to avoid the LockTimeoutException.

The following code shows how to create an ObjectGrid object, define a map, and set its LockTimeout value to 30 seconds:

```
ObjectGrid objGrid = new ObjectGrid();
BackingMap bMap = objGrid.defineMap("MapName");
bMap.setLockTimeout(30);
```

Use the previous hardcoded example to set ObjectGrid and map properties. If you create ObjectGrid from an XML file, set the **LockTimeout** attribute within the backingMap element. The following is an example of a backingMap element that sets a map LockTimeout value to 30 seconds.

```
<backingMap name="MapName" lockStrategy="PESSIMISTIC" lockTimeout="30">
```

- **Problem:** Single key deadlocks.

Description: The following scenarios describe how deadlocks can occur when a single key is accessed using a S lock and later updated. When this happens from two transactions simultaneously, it results in a deadlock.

Table 9. Single key deadlocks scenario

| | Thread 1 | Thread 2 | |
|---|--------------------|--------------------|--|
| 1 | session.begin() | session.begin() | Each thread establishes an independent transaction. |
| 2 | map.get(key1) | map.get(key1) | S lock granted to both transactions for key1. |
| 3 | map.update(Key1,v) | | No U lock. Update performed in transactional cache. |
| 4 | | map.update(key1,v) | No U lock. Update performed in the transactional cache |
| 5 | session.commit() | | Blocked: The S lock for key1 cannot be upgraded to an X lock because Thread 2 has an S lock. |
| 6 | | session.commit() | Deadlock: The S lock for key1 cannot be upgraded to an X lock because T1 has an S lock. |

Table 10. Single key deadlocks, continued

| | Thread 1 | Thread 2 | |
|---|--------------------------|--------------------------|---|
| 1 | session.begin() | session.begin() | Each thread establishes an independent transaction. |
| 2 | map.get(key1) | | S lock granted for key1 |
| 3 | map.getForUpdate(key1,v) | | S lock is upgraded to a U lock for key1. |
| 4 | | map.get(key1) | S lock granted for key1. |
| 5 | | map.getForUpdate(key1,v) | Blocked: T1 already has U lock. |
| 6 | session.commit() | | Deadlock: The U lock for key1 cannot be upgraded. |
| 7 | | session.commit() | Deadlock: The S lock for key1 cannot be upgraded. |

Table 11. Single key deadlocks, continued

| | Thread 1 | Thread 2 | |
|---|--------------------------|--------------------------|--|
| 1 | session.begin() | session.begin() | Each thread establish an independent transaction |
| 2 | map.get(key1) | | S lock granted for key1. |
| 3 | map.getForUpdate(key1,v) | | S lock is upgraded to a U lock for key1 |
| 4 | | map.get(key1) | S lock is granted for key1. |
| 5 | | map.getForUpdate(key1,v) | Blocked: Thread 1 already has a U lock. |

Table 11. Single key deadlocks, continued (continued)

| | Thread 1 | Thread 2 | |
|---|------------------|----------|---|
| 6 | session.commit() | | Deadlock: The U lock for key1 cannot be upgraded to an X lock because Thread 2 has an S lock. |

If the `ObjectMap.getForUpdate` is used to avoid the S lock, then the deadlock is avoided:

Table 12. Single key deadlocks, continued

| | Thread 1 | Thread 2 | |
|---|------------------------|------------------------|--|
| 1 | session.begin() | session.begin() | Each thread establishes an independent transaction. |
| 2 | map.getForUpdate(key1) | | U lock granted to thread 1 for key1. |
| 3 | | map.getForUpdate(key1) | U lock request is blocked. |
| 4 | map.update(key1,v) | <blocked> | |
| 5 | session.commit() | <blocked> | The U lock for key1 can be successfully upgraded to an X lock. |
| 6 | | <released> | The U lock is finally granted to key1 for thread 2. |
| 7 | | map.update(key2,v) | U lock granted to thread 2 for key2. |
| 8 | | session.commit() | The U lock for key1 can successfully be upgraded to an X lock. |

Solutions:

1. Use the `getForUpdate` method instead of `get` to acquire a U lock instead of an S lock.
 2. Use a transaction isolation level of read committed to avoid holding S locks. Reducing the transaction isolation level increases the possibility of non-repeatable reads. However, non-repeatable reads from one client are only possible if the transaction cache is explicitly invalidated by the same client.
 3. Use the optimistic lock strategy. Using the optimistic lock strategy requires handling optimistic collision exceptions.
- **Problem:** Ordered multiple key deadlock

Description: This scenario describes what happens if two transactions attempt to update the same entry directly and hold S locks to other entries.

Table 13. Ordered multiple key deadlock scenario

| | Thread 1 | Thread 2 | |
|---|--------------------|--------------------|---|
| 1 | session.begin() | session.begin() | Each thread establishes an independent transaction. |
| 2 | map.get(key1) | map.get(key1) | S lock granted to both transactions for key1. |
| 3 | map.get(key2) | map.get(key2) | S lock granted to both transactions for key2. |
| 4 | map.update(key1,v) | | No U lock. Update performed in transactional cache. |
| 5 | | map.update(key2,v) | No U lock. Update performed in transactional cache. |

Table 13. Ordered multiple key deadlock scenario (continued)

| | Thread 1 | Thread 2 | |
|----|------------------|------------------|---|
| 6. | session.commit() | | Blocked: The S lock for key 1 cannot be upgraded to an X lock because thread 2 has an S lock. |
| 7 | | session.commit() | Deadlock: The S lock for key 2 cannot be upgraded because thread 1 has an S lock. |

You can use the `ObjectMap.getForUpdate` method to avoid the S lock, then you can avoid the deadlock:

Table 14. Ordered multiple key deadlock scenario, continued

| | Thread 1 | Thread 2 | |
|----|------------------------|------------------------|--|
| 1 | session.begin() | session.begin() | Each thread establishes an independent transaction. |
| 2 | map.getForUpdate(key1) | | U lock granted to transaction T1 for key1. |
| 3 | | map.getForUpdate(key1) | U lock request is blocked. |
| 4 | map.get(key2) | <blocked> | S lock granted for T1 for key2. |
| 5 | map.update(key1,v) | <blocked> | |
| 6 | session.commit() | <blocked> | The U lock for key1 can be successfully upgraded to an X lock. |
| 7 | | <released> | The U lock is finally granted to key1 for T2 |
| 8 | | map.get(key2) | S lock granted to T2 for key2. |
| 9 | | map.update(key2,v) | U lock granted to T2 for key2. |
| 10 | | session.commit() | The U lock for key1 can be successfully upgraded to an X lock. |

Solutions:

1. Use the `getForUpdate` method instead of the `get` method to acquire a U lock directly for the first key. This strategy works only if the method order is deterministic.
2. Use a transaction isolation level of read committed to avoid holding S locks. This solution is the easiest to implement if the method order is not deterministic. Reducing the transaction isolation level increases the possibility of non-repeatable reads. However, non-repeatable reads are only possible if the transaction cache is explicitly invalidated.
3. Use the optimistic lock strategy. Using the optimistic lock strategy requires handling optimistic collision exceptions.

- **Problem:** Out of order with U lock

Description: If the order in which keys are requested cannot be guaranteed, then a deadlock can still occur.

Table 15. Out of order with U lock scenario

| | Thread 1 | Thread 2 | |
|---|------------------------|------------------------|---|
| 1 | session.begin() | session.begin() | Each thread establishes an independent transaction. |
| 2 | map.getforUpdate(key1) | map.getForUpdate(key2) | U locks successfully granted for key1 and key2. |

Table 15. Out of order with U lock scenario (continued)

| | Thread 1 | Thread 2 | |
|---|--------------------|--------------------|--|
| 3 | map.get(key2) | map.get(key1) | S lock granted for key1 and key2. |
| 4 | map.update(key1,v) | map.update(key2,v) | |
| 5 | session.commit() | | The U lock cannot be upgraded to an X lock because T2 has an S lock. |
| 6 | | session.commit() | The U lock cannot be upgraded to an X lock because T1 has an S lock. |

Solutions:

1. Wrap all work with a single global U lock (mutex). This method reduces concurrency, but handles all scenarios when access and order is non-deterministic.
2. Use a transaction isolation level of read committed to avoid holding S locks. This solution is the easiest to implement if the method order is not deterministic and provides the greatest amount of concurrency. Reducing the transaction isolation level increases the possibility of non-repeatable reads. However, non-repeatable reads are only possible if the transaction cache is explicitly invalidated.
3. Use the optimistic lock strategy. Using the optimistic lock strategy requires handling optimistic collision exceptions.

Troubleshooting security

Use this information to troubleshoot issues with your security configuration.

Procedure

- **Problem:** The client end of the connection requires Secure Sockets Layer (SSL), with the transportType setting set to SSL-Required. However, the server end of the connection does not support SSL, and has the transportType setting set to TCP/IP. As a result, the following exception gets chained to another exception in the log files:

```
java.net.ConnectException: connect: Address is invalid on local machine, or
port is not valid on remote machine
    at java.net.PlainSocketImpl.doConnect(PlainSocketImpl.java:389)
    at java.net.PlainSocketImpl.connectToAddress(PlainSocketImpl.java:250)
    at java.net.PlainSocketImpl.connect(PlainSocketImpl.java:237)
    at java.net.SocksSocketImpl.connect(SocksSocketImpl.java:385)
    at java.net.Socket.connect(Socket.java:540)
    at
com.ibm.rmi.transport.TCPTransportConnection.createSocket(TCPTransportConnection.java:155)
    at
com.ibm.rmi.transport.TCPTransportConnection.createSocket(TCPTransportConnection.java:167)
```

The address in this exception could be a catalog server, container server, or client.

Solution: See the information on configuring secure transport types in the *Administration Guide* for a table with the valid security configurations between clients and servers.

- When agent is used, the client sends the agent call to the server, and server sends the response back to the client to acknowledge the agent call. When the agent finishes processing, the server will initiate a connect to send the agent

results. This makes the container server a client from connect point of view. Therefore, if TLS/SSL is configured, make sure the client's public certificate is imported in the server's trust store.

IBM Support Assistant for WebSphere eXtreme Scale

You can use the IBM Support Assistant to collect data, analyze symptoms, and access product information.

IBM Support Assistant Lite

IBM Support Assistant Lite for WebSphere eXtreme Scale provides automatic data collection and symptom analysis support for problem determination scenarios.

IBM Support Assistant Lite reduces the amount of time it takes to reproduce a problem with the proper Reliability, Availability, and Serviceability tracing levels set (trace levels are set automatically by the tool) to streamline problem determination. If you need further assistance, IBM Support Assistant Lite also reduces the effort required to send the appropriate log information to IBM Support.

IBM Support Assistant Lite is included in each installation of WebSphere eXtreme Scale Version 7.1.0

IBM Support Assistant

IBM® Support Assistant (ISA) provides quick access to product, education, and support resources that can help you answer questions and resolve problems with IBM software products on your own, without needing to contact IBM Support. Different product-specific plug-ins let you customize IBM Support Assistant for the particular products you have installed. IBM Support Assistant can also collect system data, log files, and other information to help IBM Support determine the cause of a particular problem.

IBM Support Assistant is a utility to be installed on your workstation, not directly onto the WebSphere eXtreme Scale server system itself. The memory and resource requirements for the Assistant could negatively affect the performance of the WebSphere eXtreme Scale server system. The included portable diagnostic components are designed for minimal impact to the normal operation of a server.

You can use IBM Support Assistant to help you in the following ways:

- To search through IBM and non-IBM knowledge and information sources across multiple IBM products to answer a question or solve a problem
- To find additional information through product-specific Web resources; including product and support home pages, customer news groups and forums, skills and training resources and information about troubleshooting and commonly asked questions
- To extend your ability to diagnose product-specific problems with targeted diagnostic tools available in the Support Assistant
- To simplify collection of diagnostic data to help you and IBM resolve your problems (collecting either general or product/symptom-specific data)
- To help in reporting of problem incidents to IBM Support through a customized online interface, including the ability to attach the diagnostic data referenced above or any other information to new or existing incidents

Finally, you can use the built-in Updater facility to obtain support for additional software products and capabilities as they become available. To set up IBM Support Assistant for use with WebSphere eXtreme Scale, first install IBM Support Assistant using the files provided in the downloaded image from the IBM Support Overview Web page at: http://www-947.ibm.com/support/entry/portal/Overview/Software/Other_Software/IBM_Support_Assistant. Next, use IBM Support Assistant to locate and install any product updates. You can also choose to install plug-ins available for other IBM software in your environment. More information and the latest version of the IBM Support Assistant are available from the IBM Support Assistant Web page at: <http://www.ibm.com/software/support/isa/>.

Messages

When you encounter a message in a log or other parts of the product interface, you can look up the message by its component prefix to find out more information.

Finding messages

When you encounter a message in a log, copy the message number with its letter prefix and number and search in the information center (for example, CWOBJ1526I). When you search for the message, you can find an additional explanation of the message and possible actions you can take to resolve the problem.

See the information center for an index of product messages.

Release notes

Links are provided to the product support Web site, to product documentation, and to last minute updates, limitations, and known problems for the product.

- “Accessing last-minute updates, limitations, and known problems”
- “Accessing system and software requirements”
- “Accessing product documentation” on page 370
- “Accessing the product support Web site” on page 370
- “Contacting IBM Software Support” on page 370

Accessing last-minute updates, limitations, and known problems

The release notes are available on the product support site as technotes. To see a list of all the technotes for WebSphere eXtreme Scale, go to the Support Web page. Clicking the links provided here will result in a search of the Support Web page for the relevant release notes, which will be returned as a list.

- **7.1+** To see a list of the release notes for Version 7.1, go to the Support Web page.
- To see a list of the release notes for Version 7.0, go to the Support Web page.
- To see a list of the release notes for Version 6.1, go to the Release notes wiki page.

Accessing system and software requirements

The hardware and software requirements are documented on the following pages:

- Detailed system requirements

Accessing product documentation

For the entire information set, go to the Library page.

Accessing the product support Web site

To search for the latest technotes, downloads, fixes, and other support-related information, go to the Support page.

Contacting IBM Software Support

If you encounter a problem with the product, first try the following actions:

- Follow the steps described in the product documentation
- Look for related documentation in the online help
- Look up error messages in the message reference

If you cannot resolve your problem by any of the preceding methods, contact IBM Technical Support.

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, New York 10594 USA

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Mail Station P300
522 South Road
Poughkeepsie, NY 12601-5400
USA
Attention: Information Requests

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

Trademarks

The following terms are trademarks of IBM Corporation in the United States, other countries, or both:

- AIX
- CICS[®]
- Cloudscape
- DB2
- Domino[®]
- IBM
- Lotus[®]
- RACF[®]
- Redbooks[®]
- Tivoli
- WebSphere
- z/OS[®]

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

LINUX is a trademark of Linus Torvalds in the U.S., other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

Index

A

- accessing data
 - ObjectGrid shard 24
 - overview 33
 - with applications 9
- Administration API 263
- API 261
- authorization 326

B

- backing map
 - lock strategy 131
 - plug-ins 12
 - session 12
- batchUpdate method 240
- best practices 207, 348
- byte array maps 41, 346

C

- clients
 - configuring 160
- connecting
 - to a distributed data grid 16
- CopyMode
 - best practices 36, 341
- create ObjectGrid 17

D

- Data 33
- data access
 - partitions 33
 - queries 33
 - stored data 33
 - transactions 33
- DataGrid API
 - example 156
 - overview 155
 - partitioning with 155
- deadlock
 - troubleshooting 362
- deadlocks
 - scenarios for 141
- distributing changes
 - using Java message service 134
- dynamic maps
 - maps 54

E

- entity
 - life cycles of 74
 - listener 79
 - schema 61
- entity manager
 - fetch plan 81
 - tutorial 92

- entity maps
 - creating 240
- entity metadata
 - emd.xsd file 68
 - XML configuration 68
- entity schema
 - entity 61
- EntityManager API
 - distributed 71
 - fetch plan 81
 - for caching objects 60
 - performance 84
 - simple queries for 105
- EntityTransaction interface 91
- event listeners 221
- evictor 165
- evictors 195
 - configuring 4, 5, 8
 - plug-in 200
 - TTL evictor 198
- exception handling 145, 154
- exclusive lock 141
- external transaction manager 257

F

- FetchPlan 81
- FIFO queues
 - maps 57

G

- get method 240
- get ObjectGrid instance 21
- grid authorization 333

H

- heaps 207, 348

I

- IBM Support Assistant 368
- index
 - callback 25
 - configuration 33
 - data access 25
 - data quality 28
 - HashIndex 33
 - non-key 25
 - performance 28
- indexing
 - composite index 227
 - hash index 227
- instrumentation agent 85
- isolation
 - for transactions 152
 - pessimistic locking 152
 - repeatable read 152

J

- Java Authentication and Authorization Service
 - JAAS 333
- Java Persistence API (JPA)
 - JPAEntityLoader plug-in
 - introduction 238
 - preload utility
 - overview 284
 - time-based data updater
 - overview 294
 - time-based updater
 - starting 291
 - using with eXtreme Scale
 - overview 281
- Java virtual machine 339
- JavaMap interface 57
- JVM 339

L

- listeners
 - for backing map objects 221
 - for the eXtreme Scale 221
 - introduction 221
 - MapEventListener plug-in 222
 - ObjectGridEventListener 223
 - ObjectGridEventListener plug-in 223
- loader 165
 - Java Persistence API (JPA)
 - overview 281
 - JPA programming considerations 236
 - overview 230
 - using with entity maps and tuples 240
 - writing 232
- locking
 - configuring programmatically 145
 - configuring with XML 145
 - no 145
 - optimistic 131, 145
 - performance 149, 350
 - pessimistic 131, 145
 - strategies for 131
- locks
 - compatibility 141
 - life cycle 141
 - time out 141
- log element 165
- log sequence 165
- LogElement 165
- LogSequence 165

M

- map entry locks
 - indexes 150
 - query 150
- messages 369

- monitoring
 - with PMI 268
 - with the statistics API 265

O

- ObjectGrid interface
 - overview 9
- ObjectGridManager interface
 - controlling life cycle with 22
 - createObjectGrid methods 17
 - getObjectGrid methods 21
 - removeObjectGrid methods 22
 - using to interact with an ObjectGrid 17
- ObjectMap API
 - caching objects with 50
 - overview 50
- ObjectTransformer
 - best practices for 215, 353

P

- partitions
 - transactions 135
 - using non-keys to find objects in 127
- performance 207, 339, 348
 - best practices 149, 350
 - locking 149, 350
- Performance Monitoring Infrastructure
 - enabling 268
 - modules 272
 - monitoring with 268
 - retrieving statistics from 270
- Performance Monitoring Infrastructure (PMI) 297
- plug-in
 - index 30, 225
 - introduction 193
 - introduction to 193
 - ObjectTransformer plug-in 210
 - OptimisticCallback 217
 - plug-in slots 254
 - TransactionCallback 250
 - WebSphereTransactionCallback plug-in 259
- PMI i
 - MBean 9, 297
- Programming eXtreme Scale 1

Q

- query 227
 - Backus Naur 114
 - BNF 114
 - clauses 106
 - client failure 87
 - entity 102
 - example 105
 - functions 106
 - get plan 117
 - index 105, 120
 - key collision 87
 - methods 92
 - object map 97
 - ObjectQuery schema 100

- query (*continued*)
 - optimization with indexes 120
 - pagination 105
 - parameters 105
 - predicates 106
 - query plan 117
 - queue 87
 - schema 100
 - search elements 92
 - tuning 116, 354
 - valid attributes 100
- queues 207, 348

R

- release notes 369
- replica preload 245
- replication
 - enabling client-side 168
- request
 - per-container 47
 - routing 47
 - Session 47

S

- security
 - local 333
 - plug-ins 333
 - programming 307
- securityAPI 307
- serialization
 - locking 214, 351
 - performance 214, 351
- session
 - collision 154
 - transaction 154
 - with ObjectGrid 9
- SessionHandle
 - routing 47
- sessions
 - access data
 - flush 43
 - grid 43
- shared lock 141
- Spring
 - extension beans 297, 300
 - framework 297
 - namespace 300
 - namespace support 297
 - native transactions 297
 - packaging 297
 - shard scope 297
 - webflow 297
- stand-alone 169
- start server
 - programmatically 263
- starting servers 169
- statistics
 - with the statistics API 265
- statistics API 297
- stop server
 - programmatically 263
- support 369
- Support 368
- system API 193

T

- time zones
 - inserting data 97
 - querying data in 95
- trace
 - options for configuring 359
- transaction 257
- transactions
 - cross-grid 135
 - overview 128, 130, 254
 - programming for 128
 - programmaming for 128
 - single-partition 135
- troubleshooting 357
 - messages 369
 - release notes 369
- TTL evictor 195
- tuning 339
- tuple objects
 - creating 240

U

- upgradeable lock 141

W

- writing evictors
 - RollBack evictor 203



Printed in USA