

Note

Before using this information, be sure to read the general information under “Notices” on page 37.

Compilation date: April 20, 2004

© Copyright International Business Machines Corporation 2004. All rights reserved.

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

Contents

How to send your comments v

Chapter 1. Welcome to Resources 1

Chapter 2. Using enterprise beans in applications 3

Enterprise beans 3

Developing enterprise beans 4

 Migrating enterprise bean code to the supported specification 5

 WebSphere extensions to the Enterprise JavaBeans specification 8

 Best practices for developing enterprise beans . . . 9

 Unknown primary-key class. 13

Using access intent policies 13

 Access intent policies 14

 Applying access intent policies to methods. . . 16

 Access intent exceptions 17

 Access intent assembly settings. 18

 Access intent best practices 20

 Frequently asked questions: Access intent . . . 20

EJB modules 21

Assembling EJB modules 22

 Container transactions. 22

 Method extensions 23

 Method permissions 23

 References. 23

EJB containers 23

Managing EJB containers 24

 EJB container settings 24

 EJB container system properties 25

 EJB cache settings 26

 Container interoperability 27

Deploying EJB modules 32

 EJB module collection 32

 EJB module settings 32

Enterprise beans: Resources for learning. 33

EJB method Invocation Queuing 34

Notices 37

Trademarks and service marks 39

How to send your comments

Your feedback is important in helping to provide the most accurate and highest quality information.

- To send comments on articles in the WebSphere Application Server Information Center
 1. Display the article in your Web browser and scroll to the end of the article.
 2. Click on the **Feedback** link at the bottom of the article, and a separate window containing an e-mail form appears.
 3. Fill out the e-mail form as instructed, and click on **Submit feedback** .
- To send comments on PDF books, you can e-mail your comments to: **wasdoc@us.ibm.com** or fax them to 919-254-0206.

Be sure to include the document name and number, the WebSphere Application Server version you are using, and, if applicable, the specific page, table, or figure number on which you are commenting.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Chapter 1. Welcome to Resources

Note: . If you would prefer to browse PDF versions of this documentation using your Adobe Reader, see the **Resources** PDF files available from www.ibm.com/software/webservers/appserv/infocenter.html.

The product supports all of the resources defined by the Java 2 Platform, Enterprise Edition (J2EE).

Data access (JDBC and J2C)

The J2EE Connector architecture defines a standard architecture that enables the integration of various enterprise information systems (EIS) with application servers and enterprise applications. It defines a standard resource adapter used by a Java application to connect to an EIS. This resource adapter can plug into the application server and, through the Common Client Interface (CCI), provide connectivity between the EIS, the application server, and the enterprise application.

For more information, refer to Welcome to Data Access.

Messaging

The product supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface. The base JMS support enables IBM WebSphere Application Server applications to exchange messages asynchronously with other JMS clients by using JMS destinations (queues or topics).

For more information, refer to Welcome to Messaging.

Mail

Using JavaMail API, a code segment can be embedded in any Java 2 Enterprise Edition (J2EE) application component, such as an EJB or a servlet, allowing the application to send a message and save a copy of the mail to the Sent folder.

For more information, refer to Welcome to Mail.

URLs

Java 2 Platform, Enterprise Edition (J2EE) applications can use URLs as resources in the same way other J2EE resources, such as JDBC and JavaMail, are used.

For more information,, refer to Welcome to URLs.

Resource environment entries

A resource environment reference maps a logical name used by the client application to the physical name of an object.

For more information, see Resource environment entries.

Chapter 2. Using enterprise beans in applications

1. Design a J2EE application and the enterprise beans that it needs. See "Resources for learning" for links to design information that is specific to enterprise beans.
2. Develop any enterprise beans that your application will use.
3. Prepare for assembly. For your EJB 2.x-compliant entity beans, decide on an appropriate access intent policy.
4. Assemble the beans using the Assembly Toolkit into one or more EJB modules. This includes setting security.
5. Assemble the modules into a J2EE application using the Assembly Toolkit .
6. **5.1+** For a given application server, update the EJB container configuration if needed for the application to be deployed, and determine if you want to batch commands batch commands or defer commands for container managed persistence.
7. Deploy the application in an application server.
8. Test the modules.
 - As needed, debug problems with the container.
 - Debug access and deployment problems.
9. Assemble the production application using the Assembly Toolkit.
10. Deploy the application to a production environment.
11. Manage the application:
 - a. Manage installed EJB modules. After an application has been installed, you can manage its EJB modules individually through administrative console settings.
 - b. Manage other aspects of the J2EE application.
12. Update the module and redeploy it using the Assembly Toolkit.
13. Tune the performance of the application. See Best practices for developing enterprise beans.

Enterprise beans

An enterprise bean is a Java component that can be combined with other resources to create J2EE applications. There are three types of enterprise beans, *entity* beans, *session* beans, and *message-driven* beans.

All beans reside in EJB containers, which provide an interface between the beans and the application server on which they reside.

Entity beans store permanent data. Entity beans with container-managed persistence (CMP) require connections to a form of persistent storage. This storage might be a database, an existing legacy application, a file, or another type of persistent storage. Entity beans with bean-managed persistence manage permanent data in whichever manner is defined in the bean code. For example, they can write data to databases or XML files

Session beans do not require database access, although they can obtain it indirectly as needed through entity beans. Session beans can also obtain direct access to databases (and other resources) through the use of resource references. Session beans can be either *stateful* or *stateless*.

Message-driven beans are new in version 2.0 of the Enterprise JavaBeans (EJB) specification. They enable asynchronous message servicing. The EJB container and a Java Message Service (JMS) provider work together to process messages. When a message arrives from another application component through JMS, the EJB container forwards it through an `onMessage()` call to a message-driven bean instance, which then processes the message. In other respects, message-driven beans are similar to stateless session beans.

Beans that require data access use *data sources*, which are administrative resources that define pools of connections to persistent storage mechanisms.

For more information about enterprise beans, see "Resources for learning."

Developing enterprise beans

Design a J2EE application and the enterprise beans that it needs.

- For general design information, see "Resources for learning."
- Before developing entity beans with container-managed persistence (CMP), read "Concurrency control."

There are two basic approaches to selecting tools for developing enterprise beans:

- You can use one of the available integrated development environments (IDEs). IDE tools automatically generate significant parts of the enterprise bean code and contain integrated tools for packaging and testing enterprise beans. The IBM WebSphere Application Developer product is the recommended IDE. For more information, see the documentation for that product.
- If you have decided to develop enterprise beans without an IDE, you need at least an ASCII text editor. You can also use a Java development tool that does not support enterprise bean development. You can then use tools available in the Java Software Development Kit (SDK) and in this product to assemble, test, and deploy the beans.

The following steps primarily support the second approach, development without an IDE.

1. If necessary, migrate any pre-existing code to the required version of the Enterprise JavaBeans (EJB) specification.
2. Write and compile the components of the enterprise bean.
 - At a minimum, an EJB 1.1 session bean requires a bean class, a home interface, and a remote interface. An EJB 1.1 entity bean requires a bean class, a primary-key class, a home interface, and a remote interface.
 - At a minimum, an EJB 2.0 session bean requires a bean class, a home or local home interface, and a remote or local interface. An EJB 2.0 entity bean requires a bean class, a primary-key class, a remote home or local home interface, and a remote or local interface. The types of interfaces go together: If you implement a local interface, you must define a local home interface as well.

Note: Optionally, the primary-key class can be *unknown*. See unknown primary-key class for more information.

- Available only through EJB 2.0, a message-driven bean requires only a bean class.
3. For each entity bean, complete work to handle persistence operations.
 - Create a database schema for the entity bean's persistent data.
 - For entity beans with container-managed persistence (CMP), you must store the bean's persistent data in one of the supported databases.

WebSphere Application Server application assembly tools automatically generate SQL code for creating database tables for CMP entity beans. If your CMP beans require complex database mappings, it is recommended that you use the IBM WebSphere Studio Application Developer product to generate code for the database tables.

- For entity beans with bean-managed persistence (BMP), you can create the database and database table by using the database tools or use an existing database and database table.

For more information on creating databases and database tables, consult your database documentation.

- **(CMP entity beans for EJB 2.0 only)** Define finder queries with EJB Query Language (EJB QL).

With EJB QL, you define finders in terms of CMP fields and container-managed relationships, as follows:

- *Public* finders are visible in the bean's home interface. Implemented in the bean class, they return only remote interfaces and collection types.
- *Private* finders, expressed as SELECT statements, are used only within the bean class. They can return both local and remote interfaces, dependent values, other CMP field types, and collection types.

- **(CMP entity beans for EJB 1.1 only: an IBM extension)** Create a finder helper interface for each CMP entity bean that contains specialized finder methods (other than the `findByPrimaryKey` method).

The following logic is required for each finder method (other than the `findByPrimaryKey` method) contained in the home interface of an entity bean with CMP:

- The logic must be defined in a public interface named *NameBeanFinderHelper*, where *Name* is the name of the enterprise bean (for example, *AccountBeanFinderHelper*).
- The logic must be contained in a String constant named *findMethodName* where *Clause*, where *findMethodName* is the name of the finder method. The String constant can contain zero or more question marks (?) that are replaced from left to right with the value of the finder method's arguments when that method is called.

5.1+ Assemble the beans in one or more EJB modules.

Migrating enterprise bean code to the supported specification

Support for Version 2.0 of the Enterprise JavaBeans (EJB) specification is new for Version 5 of this product. Migration of enterprise beans deployed in Version 4.0.x of this product is not generally necessary; Version 1.1 of the EJB specification is still supported. Follow these steps as appropriate for your application deployment.

1. Modify enterprise bean code for changes in the specification.
 - For Version 1.0 beans, migrate at least to Version 1.1.
 - As stated previously, migration from Version 1.1 to Version 2.0 of the EJB specification is not required for redeployment on this version of the product. However, if your application requires the capabilities of Version 2.0, migrate your Version 1.1-compliant code.

Note: The EJB Version 2.0 specification mandates that prior to the EJB container's executing a `findByMethod` query, the state of all enterprise beans enlisted in the current transaction be synchronized with the persistent store. (This is so the query is performed against current data.) If Version 1.1 beans are reassembled into an EJB 2.0-compliant

module, the EJB container synchronizes the state of Version 1.1 beans as well as that of Version 2.0 beans. As a result, you might notice some change in application behavior even though the application code for the Version 1.1 beans has not been changed.

2. Modify enterprise bean code for changes in deployment requirements. If the enterprise beans were previously deployed in Version 3.0.x of this product, modify import statements to match standard package names. In Version 3.0.2.x, the following standard packages were present under nonstandard names:

```
javax.sql.*  
javax.transaction.*
```

Any code using WebSphere data sources, including BMP entity beans and session beans that access databases, must be modified.

3. You might have to modify code for some EJB 1.1-compliant modules that were not migrated to Version 2.0. Use the following information to help you decide.
 - Some stub classes for deployed enterprise beans have changed in the Java 2 Software Development Kit, Version 1.4.1.
 - The task of generating deployment code for enterprise beans changed significantly for EJB 1.1-compliant modules relative to EJB 1.0-compliant modules.
 - If the CMP beans write to databases with mixed-case table or column names and you used IBM VisualAge for Java, Version 3.5.x, to generate the original deployment code, you cannot simply reassemble the beans in this product. You must export the original EJB project from the VisualAge for Java product as an EJB 1.1 JAR. This preserves the metadata needed to generate the correct deployment code for mixed-case database tables and columns. For more information, see the documentation for the Deployment Tool for Enterprise JavaBeans.

For detailed information about source and binary compatibility between deployed versions, see "Resources for learning."

4. Reassemble and redeploy all modules to incorporate migrated code.

Migrating enterprise bean code from Version 1.0 to Version 1.1

The following information generally applies to any enterprise bean that currently complies with Version 1.0 of the Enterprise JavaBeans (EJB) specification. For more information about migrating code for beans produced with the IBM WebSphere Studio Application Developer tool, see the documentation for that product. For more information about migrating code in general, see "Resources for learning."

1. In session beans, replace all uses of `javax.jts.UserTransaction` with `javax.transaction.UserTransaction`. Entity beans may no longer use the `UserTransaction` interface at all.
2. In finder methods for entity beans, include `FinderException` in the throws clause.
3. Remove throws of `java.rmi.RemoteException`; throw `javax.ejb.EJBException` instead. However, continue to include `RemoteException` in the throws clause of home and remote interfaces as required by the use of Remote Method Invocation (RMI).
4. Remove uses of the `finalize()` method.
5. Replace calls to `getCallerIdentity()` with calls to `getCallerPrincipal()`. The use of `getCallerIdentity()` is deprecated.
6. Replace calls to `isCallerInRole(Identity)` with calls to `isCallerInRole(String)`. The use of `isCallerInRole(Identity)` and `java.security.Identity` is deprecated.

7. Replace calls to `getEnvironment()` in favor of JNDI lookup. As an example, change the following code:

```
String homeName =
    aLink.getEntityContext().getEnvironment().getProperty("TARGET_HOME_NAME");
if (homeName == null) homeName = "TARGET_HOME_NAME";
```

The updated code would look something like the following:

```
Context env = (Context)(new InitialContext()).lookup("java:comp/env");
String homeName = (String)env.lookup("ejb10-properties/TARGET_HOME_NAME");
```

8. In `ejbCreate` methods for an entity bean with container-managed persistence (CMP), return the bean's primary key class instead of `void`.
9. Add the `getHomeHandle()` method to home interfaces.

```
public javax.ejb.HomeHandle getHomeHandle() {return null;}
```

Consider enhancements to match the following changes in the specification:

- Primary keys for entity beans can be of type `java.lang.String`.
- Finder methods for entity beans return `java.util.Collection`.
- Check the format of any JNDI names being used. Local name spaces are also supported.
- Security is defined by role, and isolation levels are defined at the method level rather than at the bean level.

Migrating enterprise bean code from Version 1.1 to Version 2.0

Enterprise JavaBeans (EJB) Version 2.0-compliant beans may be assembled only in an EJB 2.0-compliant module, although an EJB 2.0-compliant module can contain a mixture of Version 1.x and Version 2.0 beans.

The EJB Version 2.0 specification mandates that prior to the EJB container's executing a *findByMethod* query, the state of all enterprise beans enlisted in the current transaction be synchronized with the persistent store. (This is so the query is performed against current data.) If Version 1.1 beans are reassembled into an EJB 2.0-compliant module, the EJB container synchronizes the state of Version 1.1 beans as well as that of Version 2.0 beans. As a result, you might notice some change in application behavior even though the application code for the Version 1.1 beans has not been changed.

The following information generally applies to any enterprise bean that currently complies with Version 1.1 of the EJB specification. For more information about migrating code for beans produced with the IBM WebSphere Studio Application Developer tool, see the documentation for that product. For more information about migrating code in general, see "Resources for learning."

1. In beans with container-managed persistence (CMP) version 1.x, replace each CMP field with abstract get and set methods. In doing so, you must make each bean class abstract.
2. In beans with CMP version 1.x, change all occurrences of `this.field = value` to `setField(value)`.
3. In each CMP bean, create abstract get and set methods for the primary key.
4. In beans with CMP version 1.x, create an EJB Query Language statement for each finder method.
5. In finder methods for beans with CMP version 1.x, return `java.util.Collection` instead of `java.util.Enumeration`.
6. Update handling of non-application exceptions.
 - To report non-application exceptions, throw `javax.ejb.EJBException` instead of `java.rmi.RemoteException`.

- Modify rollback behavior as needed: In EJB versions 1.1 and 2.0, all non-application exceptions thrown by the bean instance result in the rollback of the transaction in which the instance is running; the instance is discarded. In EJB 1.0, the container does not roll back the transaction or discard the instance if it throws `java.rmi.RemoteException`.
7. Update rollback behavior as the result of application exceptions.
 - In EJB versions 1.1 and 2.0, an application exception does not cause the EJB container to automatically roll back a transaction.
 - In EJB Version 1.1, the container performs the rollback only if the instance has called `setRollbackOnly()` on its `EJBContext` object.
 - In EJB Version 1.0, the container is required to roll back a transaction when an application exception is passed through a transaction boundary started by the container.

WebSphere extensions to the Enterprise JavaBeans specification

This article outlines extensions to the Enterprise JavaBeans (EJB) specification that IBM provides with this product:

Inheritance in enterprise beans

In the Java language, *inheritance* is the creation of a new class from an existing class or a new interface from an existing interface. This product supports two forms of inheritance: standard class inheritance and EJB inheritance.

In standard class inheritance, the home interface, remote interface, or enterprise bean class inherits properties and methods from base classes that are not themselves enterprise bean classes or interfaces.

By contrast in enterprise bean inheritance, an enterprise bean inherits properties (such as container-managed persistence (CMP) fields and container-managed relationship (CMR) fields), methods, and method-level control descriptor attributes from another enterprise bean.

For more information, see the documentation for the IBM WebSphere Studio Application Developer product.

Optimistic concurrency control for container-managed persistence

This product supports optimistic concurrency control of data access.

Access intents for EJB persistence

This product supports the application of named data-access policies.

Performance enhancements

Through the lifetime-in-cache settings, this product provides a way for you to improve performance for beans that are only occasionally updated. For more information, see "Entity bean assembly settings."

Some enterprise beans created with the IBM WebSphere Studio Application Developer product can utilize *read-ahead* for loading a bean and its related beans in a single database operation. An entire object graph or any part of the graph can be preloaded by configuring a finder method to use read-ahead.

Assembly and deployment extensions

5.1+ This product supports IBM extensions of assembly and deployment options.

Best practices for developing enterprise beans

Use the following guidelines when designing and developing enterprise beans:

- Use a stateless session bean to act as the entry point for business logic. For more information about using session facades, see "Resources for learning."
- Entity beans should use container-managed persistence.
- In an Enterprise JavaBeans (EJB) Version 2.0 environment, use local interfaces to improve communication between enterprise beans in the same Java virtual machine.

Local calls avoid the overhead of RMI/IIOP and use pass-by-reference semantics instead of pass-by-value. For each call, the caller and callee beans share the state of arguments. EJB 2.0 beans can have both a local and remote interface but more typically have one or the other.

- For communicating with remote clients, provide remote and remote home interfaces. For communicating with local clients like servlets, entity beans, and message-driven beans, provide local and local home interfaces.

Batch commands for container managed persistence

From JDBC 2.0 on, *PreparedStatement* objects can maintain a list of commands that can be submitted together as a batch. Instead of multiple database round trips, there can be only one database round trip for all the batched persistence requests.

The WebSphere Application Server version 5.0.2 enables you to take advantage of this. You can turn this option on from the EJB CMP side. When you choose this option, the run time defers *ejbStore/ejbCreate/ejbRemove* or the equivalent database persistence requests (insert/update/delete) until they are needed. This can be at the end of the transaction, or when a flush is needed for finders related to this EJB type. When the persistence operation finally happens, run time accumulates the database requests and uses JDBC *PreparedStatement* batch operation to make a single JDBC call for multiple rows of the same operation.

Setting the run time for batched commands:

1. Open the administrative console.
2. Select **Servers**.
3. Select **Application Servers**.
4. Select the server you want to configure.
5. In the Additional Properties area, select **Process Definition**.
6. In the Additional Properties area, select **Java Virtual Machine**.
7. Update the **Generic JVM arguments** with *Dcom.ibm.ws.pm.batch=true*.

Deferred Create for container managed persistence

The specification for Enterprise Java Beans (EJB) 2.x states that for Container Managed Persistence (CMP) during the *ejbCreate*, the container can create the representation of the entity in the database immediately, or defer it to a later time.

The WebSphere Application Server version 5.0.2 enables you to take advantage of this specification. You can turn this option on from the EJB CMP side. When you choose this option, the runtime defers *ejbCreate* (or the equivalent database persistence request) until it is needed. This can be at the end of the transaction, or

when a flush is needed for finders related to this EJB type. By doing this you can reduce two round trips for the newly created entity (insert and update) to one (insert).

Setting the run time for deferred create:

1. Open the administrative console.
2. Select **Servers**.
3. Select **Application Servers**.
4. Select the server you want to configure.
5. In the Additional Properties area, select **Process Definition**.
6. In the Additional Properties area, select **Java Virtual Machine**.
7. Update the **Generic JVM arguments** with `Dcom.ibm.ws.pm.deferredcreate=true`.

Explicit invalidation in the Persistence Manager cache

Container managed persistence (CMP) entity beans can be configured as *long-lifetime* beans. A long-lifetime bean is one that is configured with *Lifetime In Cache Usage* equal to a value other than the default **OFF**. A value other than **OFF** means that data for this bean is cached beyond the end of the transaction in which the bean was obtained by a finder or other method. The *Lifetime In Cache Usage* and *Lifetime In Cache* values control the basic length of time the cached data remains valid. When the specified time runs out, the cached data is no longer valid. See the *LifetimeInCache* help sections of the Assembly Toolkit (ATK) for more details.

However, there is also an API that lets the client application code explicitly invalidate the cached data of a bean on demand, superceding the basic lifetime of the cache data as controlled by the *Lifetime In Cache Usage* and *Lifetime In Cache* settings. This is useful where an application that does not use CMP beans modifies the data that underlies a CMP bean (for example, it updates a database table to which a CMP bean is mapped). Such an application can inform WebSphere Application Server that any cached version of this bean data is **stale** and no longer matches what is in the database. The data should be invalidated (in essence, discarded). For CMP beans that cannot tolerate stale data, this is an important feature.

Because the PM Cache Invalidation mechanism does consume resources in exchange for its benefits, it is not enabled by default. To enable it refer to Setting Persistence Manager Cache Invalidation .

Example: Explicit Invalidation in the Persistence Manager Cache: Usage Scenario

The scenario of use for this feature begins with configuring one or more bean types to be long-lifetime beans (see Explicit Invalidation in the Persistence Manager Cache, and configuring the necessary Java Message Service (JMS) resources (described below). Once this is done, the server is started. The scenario continues as follows:

1. Assume that a CMP entity bean of type *Department* has been configured to be a long-lifetime bean.
2. Transaction 1 begins. Application code looks up *Department's* home and calls a finder method (such as `findByPrimaryKey("dept01")`). As this is the first finder to return *Department dept01*, a trip is made to the database to obtain the data. Transaction 1 ends.

3. Transaction 2 begins. Application code calls `findByPrimaryKey("dept01")` again. Because this is not the first finder to return *Department dept01*, we get a cache hit and no database trip is made. So far this is current WebSphere Application Server behavior for long-lifetime beans. Transaction 2 ends.
4. Another application, which does not use the *Department* CMP bean, is executed. This application might or might not be run on the WebSphere Application Server; it could be a legacy application. The application updates the database table that is mapped to the *Department* bean, altering the row for *dept01*. For example, the *budget* column in the table (mapped to a Java double CMP attribute in the *Department* bean) is changed from \$10,000.00 to \$50,000.00. This application was designed to cooperate with WebSphere Application Server. After performing the update, the application sends an invalidate request message to invalidate the *Department* bean *dept01*.
5. Transaction 3 begins. Application code looks up *Department's* home and calls a finder method (such as `findByPrimaryKey("dept01")`). Because this is the first finder after *Department dept01* is invalidated, a new database trip is made to obtain the data. Transaction 3 ends.

Persistence Manager cache invalidation API

The PM cache invalidation API is in the form of a JMS message that the client sends to a specially-named JMS topic using a connection from a specifically named JMS *TopicConnectionFactory*. The JMS message must be an *ObjectMessage* created by the client. The client code creates a *PMCacheInvalidationRequest* object that describes the bean data to invalidate. Client code places the *PMCacheInvalidationRequest* object in the *ObjectMessage* and publishes the *ObjectMessage* (for further details on the JMS objects and terms used here, please see the Java Message Service documentation).

The public class *PMCacheInvalidationRequest* is central to the API, so we include a portion of its code here for illustration purposes (if you see any differences between this illustration and the actual class, the class is to be considered correct):

```
package com.ibm.websphere.ejbpersistence;
```

```
/**
 *An instance of this class represents a request to invalidate one or more
 *CMP beans in the PMcache. When an invalidate occurs, cached data for this
 *bean is removed from the cache; the next time an application tries to find
 *this bean, a fresh copy of the bean data is obtained from the data store.
 *
 *The ability to invalidate a bean means that a CMP bean may be configured
 *as a long-lifetime bean and thus be cached across transactions for much
 *greater performance on future attempts to find this bean. Yet when some
 *outside mechanism updates the bean data, sending an invalidation request
 *will remove stale data from the PMcache so applications do not behave falsely
 *based on stale data.
 */
public class PMCacheInvalidationRequest implements Serializable {
    . . .

    /**
     * Constructor used to invalidate a single bean
     * @param beanHomeJNDIName the JNDI name of the bean home. This is the same value
     * used to look up the bean home prior to calling findByPrimaryKey, for example.
     * @param beanKey the primary key of the bean to be invalidated. The actual
     * object type must be the primary key type for this bean type.
     */
    public PMCacheInvalidationRequest(String beanHomeJNDIName, Object beanKey)
    throws IOException {
```

```

    . . .
}
/**
 * Constructor used to invalidate a Collection of beans
 * @param beanHomeJNDIName java.lang.String the JNDI name of the bean home.
 * This is the same value used to look up the bean home prior to calling
 * findByPrimaryKey, for example.
 * @param beanKeys a Collection of the primary keys of the beans to be
 * invalidated. The actual type of each object in the Collection must be the
 * primary key type for this bean type.
 */
public PMCacheInvalidationRequest(String beanHomeJNDIName, Collection beanKeys)
    throws IOException {
    . . .
}
/**
 * Constructor used to invalidate all beans of a given type
 * @param beanHomeJNDIName java.lang.String the JNDI name of the bean home.
 * This is the same value used to look up the bean home prior to calling
 * findByPrimaryKey, for example.
 */
public PMCacheInvalidationRequest(String beanHomeJNDIName) {
    . . .
}
}
}

```

If the client wants to perform the invalidation in a synchronous way, it can opt to receive an acknowledgement JMS message when the invalidation is complete. To ask for an acknowledgement (ACK) message, the client sets a *Topic* of its own choosing in the *JMSReplyTo* field of the *ObjectMessage* for the invalidation request (see JMS documentation for further details). The client then waits (using the *receive()* method of JMS) on receipt of the acknowledgement message before continuing execution.

An ACK message enables the caller to insure there is not even a brief (seconds or less) window during which PM cache data is stale. The sending of an acknowledgement for each request does, of course, take a bit more time and so is recommended to be used only when needed.

The JMS resources used to make an invalidation request (*TopicConnectionFactory*, *TopicDestination*, and so forth) must be configured by the user (using the Administration console or other method) if they want to use PM Cache Invalidation. In this way the user can chose whichever JMS provider they prefer (as long as it supports pub-sub). The names that must be used for these resources are defined as part of the API, and use names unique to the WebSphere Application Server namespace to avoid name conflict with customer JMS resources.

The following are the names that must be used when the user configures the JMS resources (shown as Java constants for clarity):

```

// The JNDI name of the TopicConnectionFactory
private static final String topicConnectionFactoryJNDIName =
    "com.ibm.websphere.ejbpersistence.InvalidatetCF";
// The JNDI name of the TopicDestination
private static final String topicDestinationJNDIName =
    "com.ibm.websphere.ejbpersistence.invalidate";
// The Topic name (part of the TopicDestination)
private static final String topicString = "com.ibm.websphere.ejbpersistence.invalidate";

```

Here are examples of how these constants can be used in client code:

```

// Look up the TopicConnectionFactory...
InitialContext ic = new InitialContext();
TopicConnectionFactory topicConnectionFactory =
    (TopicConnectionFactory) ic.lookup(topicConnectionFactoryJNDIName);
...
// Look up the Topic
Topic topic = (Topic) ic.lookup(topicDestinationJNDIName);

```

Note that JMS messages can be sent not only from J2EE application code (for example, a SessionBean or BMP entity bean method) but also from non-J2EE applications if your chosen JMS provider allows for this. For example, the IBM MQ provider, available in WebSphere Application Server as the **Embedded Messaging** feature (selectable during installation), supports the use of MQ classes (or structures in other languages) to create a topic connection and topic that are compatible with the *TopicConnectionFactory* and *TopicDestination* you configure using WebSphere Application Server Application Console.

Setting Persistence Manager Cache Invalidation:

1. Open the administrative console.
2. Select **Servers**.
3. Select **Application Servers**.
4. Select the server you want to configure.
5. In the Additional Properties area, select **Process Definition**.
6. In the Additional Properties area, select **Java Virtual Machine**.
7. Update the **Generic JVM arguments** with
-Dcom.ibm.ws.ejbpersistence.cacheinvalidation=true.

Unknown primary-key class

When writing an entity bean for Enterprise Java Bean Version 2.0, the minimum requirements usually include a primary-key class. However, in some cases you might choose not to specify the primary-key class for an entity bean with container managed persistence (CMP). Perhaps there is no obvious primary key, or you want to allow the deployer to select the primary key fields at deployment time. The primary key type is usually derived from the type used by the database system that stores the entity objects, and you might not know what this key is.

So, the *unknown key type* is actually a type chosen at deployment time, making it changeable each time the bean is deployed. Your client code must deal with this key as type *Object*.

Currently, WebSphere Application Server supports top-down mapping and enables the deployer to choose *String* keys generated at the application server. For an example of how to use this function, see the Samples library.

Using access intent policies

You can use access intent policies to help the product run-time environment manage various aspects of Enterprise JavaBeans (EJB) persistence. You apply access intent policies to EJB Version 2.0 entity beans and their methods by using an application assembly tool. A set of default access intent policies comes with the Assembly Toolkit.

1. Apply default access intent to CMP entity beans. For more information, see the online help available with the Assembly Toolkit.
2. Apply access intent policies to methods of CMP entity beans.

Access intent policies

An access intent policy is a named set of properties (access intents) that governs data access for Enterprise JavaBeans (EJB) persistence. You can assign policies to an entity bean and to individual methods on an entity bean's home, remote, or local interfaces during assembly. You can set access intents only within EJB Version 2.x-compliant modules for entity beans with CMP Version 2.x.

This product supplies a number of access intent policies that specify permutations of read intent and concurrency control; the pessimistic/update policy can be qualified further. The selected policy determines the appropriate isolation level and locking strategy used by the run time environment.

Access intent policies are specifically designed to supplant the use of isolation level and access intent method-level modifiers found in the extended deployment descriptor for EJB version 1.1 enterprise beans. You cannot specify isolation level and read-only modifiers for EJB version 2.0 enterprise beans.

Access intent policies configured on an entity basis define the default access intent for that entity. The default access intent controls the entity unless you specify a different access intent policy based on either method-level configuration or application profiling

You can use application profiling or method level access intent policies to control access intent more precisely. Application profiling is only available in the Integration Server product. Method-level access intent policies are named and defined at the module level. A module can have one or many such policies. Policies are assigned, and apply, to individual methods of the declared interfaces of entity beans and their associated home interfaces. A method-based policy is acted upon by the combination of the EJB container and persistence manager when the method causes the entity to load.

For entity beans that are backed by tables with nullable columns, use an optimistic policy with caution. Nullable columns are automatically excluded from overqualified updates at deployment time; concurrent changes to a nullable field might result in lost updates. When used with the IBM WebSphere Studio Application Developer product, this product provides support for selecting a subset of the non-nullable columns that are to be reflected in the overqualified update statement that is generated in the deployment code to support optimistic policies.

An entity that is configured with a read-only policy that causes a bean to be activated can cause problems if updates are attempted within the same transaction. Those changes are not committed, and the process throws an exception because data integrity might be compromised.

Concurrency control

Concurrency control is the management of contention for data resources. A concurrency control scheme is considered *pessimistic* when it locks a given resource early in the data-access transaction and does not release it until the transaction is closed. A concurrency control scheme is considered *optimistic* when locks are acquired and released over a very short period of time at the end of a transaction.

The objective of optimistic concurrency is to minimize the time over which a given resource would be unavailable for use by other transactions. This is especially important with long-running transactions, which under a pessimistic scheme would lock up a resource for unacceptably long periods of time.

Under an optimistic scheme, locks are obtained immediately before a read operation and released immediately afterwards. Update locks are obtained immediately before an update operation and held until the end of the transaction.

To enable optimistic concurrency, this product uses an *overqualified update scheme* to test whether the underlying data source has been updated by another transaction since the beginning of the current transaction. With this scheme, the columns marked for update and their original values are added explicitly through a WHERE clause in the UPDATE statement so that the statement fails if the underlying column values have been changed. As a result, this scheme can provide column-level concurrency control; pessimistic schemes can control concurrency at the row level only.

Optimistic schemes typically perform this type of test only at the end of a transaction. If the underlying columns have not been updated since the beginning of the transaction, pending updates to container-managed persistence fields are committed and the locks are released. If locks cannot be acquired or if some other transaction has updated the columns since the beginning of the current transaction, the transaction is rolled back: All work performed within the transaction is lost.

Pessimistic and optimistic concurrency schemes require different transaction isolation levels. Enterprise beans that participate in the same transaction and require different concurrency control schemes cannot operate on the same underlying data connection.

Whether or not to use optimistic concurrency depends on the type of transaction. Transactions with a high penalty for failure might be better managed with a pessimistic scheme. (A high-penalty transaction is one for which recovery would be risky or resource-intensive.) For low-penalty transactions, it is often worth the risk of failure to gain efficiency through the use of an optimistic scheme. In general, optimistic concurrency is more efficient when update collisions are expected to be infrequent; pessimistic concurrency is more efficient when update collisions are expected to occur often.

Read-ahead hints

Read-ahead schemes enable applications to minimize the number of database roundtrips by retrieving a working set of container-managed persistence (CMP) beans for the transaction within one query. Read-ahead involves activating the requested CMP beans and caching the data for their related beans, which ensures that data is present for the beans that are most likely to be needed next by an application. A *read-ahead hint* is a canonical representation of the related beans that are to be read. It is associated with the *findByPrimaryKey* method for the requested bean type, which must be an EJB 2.x-compliant CMP entity bean.

Read-ahead hints can be set only using the WebSphere Business Integration Server Foundation assembly tool or through the Add Access Intent wizard of the IBM WebSphere Studio Application Developer product.

Read-ahead is only supported for access intent policies that can be applied by the backend against which the application is deployed. Otherwise, the read-ahead hint is disregarded.

Currently, only *findByPrimaryKey* methods can have read-ahead hints. Only beans related to the requested beans by a container-managed relationship (CMR), either directly or indirectly through other beans, can be read ahead. Beans that use EJB inheritance should not be used in a read-ahead hint.

A read-ahead hint takes the form of a character string. You do not have to provide the string; the wizard generates it for you based on CMRs defined for the bean. The following example is provided as supplemental information only.

Suppose a CMP bean type A has a finder method that returns instances of bean A. A read-ahead hint for this method is specified using the following notation:
RelB.RelC; RelD

Interpret the preceding notation as follows:

- Bean type A has a CMR with bean types B and D.
- Bean type B has a CMR with bean type C.

For each bean of type A that is retrieved from the database, its directly-related B and D beans and its indirectly-related C beans are also retrieved. The order of the retrieved bean data columns in each row of the result set is the same as their order in the read-ahead hint: an A bean, a B bean (or null), a C bean (or null), a D bean (or null). For hints in which the same relationship is mentioned more than once (for example, *RelB.RelC;RelB.RelE*), a bean's data columns appear only once, at the position it first appears in the hint.

The tokens shown in the notation (*RelB* and so on) must be CMR field names for the relationships as defined in the deployment descriptor for the bean. In indirect relationships such as *RelB.RelC*, *RelC* is a CMR field name defined in the deployment descriptor for bean type B.

A single read-ahead hint cannot refer to the same bean type in more than one relationship. For example, if a Department bean has a relationship *employees* with the Employee bean and also has a relationship *manager* with the Employee bean, the read-ahead hint cannot specify both *employees* and *manager*.

For more information about how to set read-ahead hints, see the documentation for the WebSphere Studio Application Developer product.

Applying access intent policies to methods

You apply an access intent policy to a method, or set of methods, in an application's entity beans through the Assembly Toolkit.

1. Start the Assembly Toolkit.
2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. Select the **Access** tab.
5. On the right side of the **Access Intent for Entities 2.x (Method Level)** panel, select **Add**. The **Add Access Intent** panel displays.
6. Specify the **Name** for your new intent policy.
7. Select the **Access intent name** from the drop-down list.
8. Enter a **Description** to help you remember what this policy does.
9. Optional: Select **Read Ahead Hint**.
10. Click **Next**. The next **Add Access Intent** panel displays, with optional attributes.

11. Optional: Decide whether or not to overwrite these optional access intent attributes. Click on those you want to change.
12. Click **Next**. The next **Add Access Intent** panel, with a list of Enterprise Beans, displays.
13. Select one or more Enterprise Beans from the list.

Note: If you selected **Read Ahead Hint** in an earlier step, you can only select **ONE** bean at this step.

14. Click **Next**. The next **Add Access Intent** panel, with a list of methods, displays.
15. Select the methods you want to use.
16. If you *DID NOT* select **Read Ahead Hint** in an earlier step, click **Finish**. If you *DID* select the Read Ahead Hint option, you can click **Next** to specify your Read Ahead Hint for the specified bean. The next **Add Access Intent** panel, with a list of EJB preload paths, displays.
17. Edit the EJB preload path by selecting relationship roles from the **Relationship roles:** window.
18. Click **Finish**. A new entry is created in the **Access Intent for Entities 2.x (Method Level)** panel

Access intent exceptions

The following exceptions are thrown in response to the application of access intent policies:

com.ibm.ws.ejbpersistence.utilpm.PersistenceManagerException

If the method that drives the `ejbLoad()` method is configured to be read-only but updates are then made within the transaction that loaded the bean's state, an exception is thrown during invocation of the `ejbStore()` method, and the transaction is rolled back. Likewise, the `ejbRemove()` method cannot succeed in a transaction that is set as read-only. If an update hint is applied to methods of entity beans with bean-managed persistence, the same behavior and exception results. The forwarded exception object contains the message string `PMGR1103E: update instance level read only bean beanName`

This exception is also thrown if the applied access intent policy cannot be honored because a finder, `ejbSelect`, or container-managed relationship (CMR) accessor method returns an inherently read-only result. The forwarded exception object contains the message string `PMGR1001: No such DataAccessSpec - methodName`

The most common occurrence of this error is when a custom finder that contains a read-only EJB Query Language (EJB QL) statement is called with an applied access intent of `wsPessimisticUpdate` or `wsPessimisticUpdate-Exclusive`. These policies require the use of a FOR UPDATE clause on the SQL SELECT statement to be executed, but a read-only query cannot support FOR UPDATE. Other examples of read-only queries include joins; the use of ORDER BY, GROUP BY, and DISTINCT keywords.

To eliminate the exception, edit the EJB query so that it does not return an inherently read-only result or change the access intent policy being applied.

- If an update access is required, change the applied access intent setting to `wsPessimisticUpdate-WeakestLockAtLoad` or `wsOptimisticUpdate`.
- If update access is not truly required, use `wsPessimisticRead` or `wsOptimisticRead`.

- If connection sharing between entity beans is required, use `wsPessimisticUpdate-WeakestLockAtLoad` or `wsPessimisticRead`.

com.ibm.websphere.ejb.container.CollectionCannotBeFurtherAccessed

If a lazy collection is driven after it is no longer in scope, and beyond what has already been locally buffered, a `CollectionCannotBeFurtherAccessed` exception is thrown.

com.ibm.ws.exception.RuntimeWarning

If an application is configured incorrectly, a run-time warning exception is thrown as the application starts; startup is ended. You can validate an application's configuration by choosing the `verify` function. Some examples of misconfiguration include:

- A method configured with two different access intent policies
- A method configured with an undefined access intent policy

javax.ejb.NoSuchEntityException

If an update fails under optimistic concurrency because fields changed within another transaction between load and store requests, a `NoSuchEntityException` is raised and the commit fails.

Access intent assembly settings

Access intent policies contain data-access settings for use by the persistence manager. Default access intent policies are configured on the entity bean. Optionally, you can associate access intent policies with one or more methods.

These settings are applicable only for EJB 2.x-compliant entity beans that are packaged in EJB 2.x-compliant modules. Connection sharing between beans with bean-managed persistence and those with container-managed persistence is possible if they all use the same access intent policy.

Name

Specifies a name for a mapping between an access intent policy and one or more methods.

Description

Contains text that describes the mapping.

Methods - Name

Specifies the name of an enterprise bean method, or the asterisk character (*). The asterisk is used to denote all of the methods of an enterprise bean's remote and home interfaces.

Methods - Enterprise bean

Specifies which enterprise bean contains the methods indicated in the `Name` setting.

Methods - Type

Used to distinguish between a method with the same signature that is defined in both the home and remote interface. Use `Unspecified` if an access intent policy applies to all methods of the bean.

Data type	String
Range	Valid values are <code>Home</code> , <code>Remote</code> , <code>Local</code> , <code>LocalHome</code> or <code>Unspecified</code>

Methods - Parameters

Contains a list of fully qualified Java type names of the method parameters. This setting is used to identify a single method among multiple methods with an overloaded method name.

Applied access intent

Specifies how the container must manage data access for persistence. Configurable both as a default access intent for an entity and as part of a method-level access intent policy.

Data type	String
Default	wsPessimisticUpdate-WeakestLockAtLoad. With Oracle, this is the same as wsPessimisticUpdate.
Range	Valid settings are wsPessimisticUpdate, wsPessimisticUpdate-NoCollision, wsPessimisticUpdate-Exclusive, wsPessimisticUpdate-WeakestLockAtLoad, wsPessimisticRead, wsOptimisticUpdate, or wsOptimisticRead. Only wsPessimisticRead and wsOptimisticRead are valid when class-level caching is enabled in the EJB container.

This product supports lazy collections. For each segment of a collection, iterating through the collection (*next()*) does not trigger a remote method call to retrieve the next remote reference. Two policies (wsPessimisticUpdate and wsPessimisticUpdate-Exclusive) are extremely lazy; the collection increment size is set to 1 to avoid overlocking the application. The other policies have a collection increment size of 25.

If an entity is not configured with an access intent policy, the run-time environment typically uses wsPessimisticUpdate-WeakestLockAtLoad by default. If, however, the **Lifetime in cache** property is set on the bean, the default value of **Applied access intent** is wsOptimisticRead; updates are not permitted.

Additional information about valid settings follows:

Profile name	Concurrency control	Access type	Transaction isolation
wsPessimisticRead (Note 1)	pessimistic	read	For Oracle, read committed. Otherwise, repeatable read
wsPessimisticUpdate (Note 2)	pessimistic	update	For Oracle, read committed. Otherwise, repeatable read
wsPessimisticUpdate-Exclusive (Note 3)	pessimistic	update	serializable
wsPessimisticUpdate-NoCollision (Note 4)	pessimistic	update	read committed
wsPessimisticUpdate-WeakestLockAtLoad (Note 5)	pessimistic	update	Repeatable read
wsOptimisticRead	optimistic	read	read committed

Profile name	Concurrency control	Access type	Transaction isolation
wsOptimisticUpdate (Note 6)	optimistic	update	read committed
Notes: <ol style="list-style-type: none"> 1. Read locks are held for the duration of the transaction. 2. The generated SELECT FOR UPDATE query grabs locks at the beginning of the transaction. 3. SELECT FOR UPDATE is generated; locks are held for the duration of the transaction. 4. A plain SELECT query is generated. No locks are held, but updates are permitted. Use cautiously. This intent enables execution without concurrency control. 5. Where supported by the backend, the generated SELECT query does not include FOR UPDATE; locks are escalated by the persistent store at storage time if updates were made. Otherwise, the same as wsPessimisticUpdate. 6. Generated overqualified-update query forces failure if CMP column values have changed since the beginning of the transaction. <p>Be sure to review the rules for forming overqualified-update query predicates. Certain column types (for example, BLOB) are ineligible for inclusion in the overqualified-update query predicate and might affect your design.</p>			

Access intent best practices

This topic outlines issues to consider when applying access intent policies to Enterprise JavaBeans (EJB) methods.

- **Start by configuring the default access intent policy for an entity.** After your application is built and running, you can more finely tune certain access paths in your application using application profiling or method-level access intent.
- **Don't mix access types.** Avoid using both pessimistic and optimistic policies in the same transaction. For most databases, pessimistic and optimistic policies use different isolation levels. This can result in multiple database connections, which prevents you from taking advantage of the performance benefits possible through connection sharing.
- **Take care when applying wsPessimisticUpdate-NoCollision.** This policy does not ensure data integrity. No database locks are held, so concurrent transactions can overwrite each other's updates. Use this policy only if you can be sure that only one transaction will attempt to update persistent store at any given time.

Frequently asked questions: Access intent

I have not applied any access intent policies at all. My application runs just fine with a DB2 database, but it fails with an Oracle database with the following message: com.ibm.ws.ejbpersistence.utilpm.PersistanceManagerException: PMGR1001E: No such DataAccessSpec :FindAllCustomers. The backend datastore does not support the SQLStatement needed by this AccessIntent: (pessimistic update-weakestLockAtLoad)(collections: transaction/25) (resource manager prefetch: 0) (AccessIntentImpl@d23690a). Why?

If you have not configured access intent, all of your data is accessed under the default access intent policy (wsPessimisticUpdate-WeakestLockAtLoad). On DB2 databases, the weakest lock is a shared one, and the query runs without a FOR UPDATE clause. On Oracle databases, however, the weakest lock is an update lock; this means that the SQL query must contain a FOR UPDATE clause. However, not every SQL statement necessarily supports FOR UPDATE; for example, if the query is being run against multiple tables in a join, FOR UPDATE is not supported.

To avoid this problem, try either of the following:

- Modify your SQL query or reconfigure your application so that an update lock is supported

- Apply an access intent policy that supports optimistic concurrency

I am calling a finder method and I get an InconsistentAccessIntentException at run time. Why?

This can occur when you use method-level access intent policies to apply more control over how a bean instance is loaded. This exception indicates that the entity bean was previously loaded in the same transaction. This could happen if you called a multifinder method that returned the bean instance with access intent policy X applied; you are now trying to load the second bean again by calling its `findByPrimaryKey` method with access intent Y applied. Both methods must have the same access intent policy applied.

Likewise, if the entity was loaded once in the transaction using an access intent policy configured on a finder, you might have called a container-managed relationship (CMR) accessor method that returned the entity bean configured to load using that entity's default access intent.

To avoid this problem, ensure that your code does not load the same bean instance twice within the same transaction with different access intent policies applied. Avoid the use of method-level access intent unless absolutely necessary.

I have two beans in a container-managed relationship. I call `findByPrimaryKey()` on the first bean and then call `getBean2()`, a CMR accessor method, on the returned instance. At that point, I get an InconsistentAccessIntentException.

Why? You are probably using read-ahead. When you loaded the first bean, you caused the second bean to be loaded under the access intent policy applied to the finder method for the first bean. However, you have configured your CMR accessor method from the first bean to the second with a different access intent policy. CMR accessor methods are really finder methods in disguise; the run-time environment behaves as if you were trying to change the access intent for an instance you have already read from persistent store.

To avoid this problem, beans configured in a read-ahead hint are all driven to load with the same access intent policy as the bean to which the read-ahead hint is applied.

I have a bean with a one-to-many relationship to a second bean. The first bean has a pessimistic-update intent policy applied. When I try to add an instance of the second bean to the first bean's collection, I get an UpdateCannotProceedWithIntegrityException. Why?

The second bean probably has a read intent policy applied. When you add the second bean to the first bean's collection, you are not updating the first bean's state, you are implicitly modifying the second bean's state. (The second bean contains a foreign key to the first bean, which is modified.)

To avoid this problem, ensure that both ends of the relationship have an update intent policy applied if you expect to change the relationship at run time.

EJB modules

An EJB module is used to assemble one or more enterprise beans into a single deployable unit. An EJB module is stored in a standard Java archive (JAR) file.

An EJB module contains the following:

- One or more deployable enterprise beans.

- A deployment descriptor, stored in an Extensible Markup Language (XML) file. This file declares the contents of the module, defines the structure and external dependencies of the beans in the module, and describes how the beans are to be used at run time.

You can deploy an EJB module as a stand alone application, or combine it with other EJB modules or with Web modules to create a J2EE application. An EJB module is installed and run in an enterprise bean container.

For more information about EJB modules, see "Resources for learning."

Assembling EJB modules

Assemble an Enterprise JavaBeans (EJB) module to contain enterprise beans and related code artifacts. Group Web components, client code, and resource adapter code in separate modules. After assembling an EJB module, you can install it as a stand-alone application or combine it with other modules into an enterprise application.

To increase performance, break container-managed persistence (CMP) enterprise beans into several enterprise bean modules during assembly. The load time for hundreds of beans is improved by distributing the beans across several JAR files and packaging them to an EAR file. Load time is faster when the administrative server attempts to start the beans, for example, 8-10 minutes versus more than one hour when one JAR file is used.

Use the Assembly Toolkit to assemble an EJB module in any of the following ways:

- Import an existing EJB module (EJB JAR file).
 - Create a new EJB module.
 - Copy code artifacts (such as entity beans) from one EJB module into a new EJB module.
1. Start the Assembly Toolkit.
 2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
 3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
 4. Migrate enterprise bean (JAR) files created with the Application Assembly Tool (AAT) or a different tool to the Assembly Toolkit. To migrate files, import your enterprise bean files to the Assembly Toolkit.
 5. Create a new EJB module.
 6. Copy code artifacts (such as entity beans) from one EJB module into a new EJB module.
 7. Verify the contents of the new EJB module in either of the following ways:
 - In the J2EE Hierarchy view, expand **EJB Modules** and view the new module.
 - Click **Window > Show View > Navigator** to see the associated files for the EJB module in a Navigator view.

Container transactions

Container transaction properties specify how an EJB container is to manage transaction scopes for the enterprise bean's method invocations. A transaction attribute is mapped to one or more methods.

Method extensions

Method extensions are IBM extensions to the standard deployment descriptors for enterprise beans.

Method extension properties are used to define transaction isolation levels for methods, to control the delegation of a principal's credentials, and to define custom finder methods.

Method permissions

A method permission is a mapping between one or more security roles and one or more methods that a member of the role can call.

References

References are logical names used to locate external resources for enterprise applications. References are defined in the application's deployment descriptor file. At deployment, the references are bound to the physical location (global JNDI name) of the resource in the target operational environment.

This product supports the following types of references:

- An EJB reference is a logical name used to locate the home interface of an enterprise bean.
- A resource reference is a logical name used to locate a connection factory object.

These objects define connections to external resources such as databases and messaging systems. The container makes references available in a JNDI naming subcontext. By convention, references are organized as follows:

- EJB references are made available in the `java:comp/env/ejb` subcontext.
- Resource references are made available as follows:
 - JDBC DataSource references are declared in the `java:comp/env/jdbc` subcontext.
 - JMS connection factories are declared in the `java:comp/env/jms` subcontext.
 - JavaMail connection factories are declared in the `java:comp/env/mail` subcontext.
 - URL connection factories are declared in the `java:comp/env/url` subcontext.

EJB containers

An Enterprise JavaBeans (EJB) container provides a run-time environment for enterprise beans within the application server. The container handles all aspects of an enterprise bean's operation within the application server and acts as an intermediary between the user-written business logic within the bean and the rest of the application server environment.

One or more EJB modules, each containing one or more enterprise beans, can be installed in a single container.

The EJB container provides many services to the enterprise bean, including the following:

- Beginning, committing, and rolling back transactions as necessary.
- Maintaining pools of enterprise bean instances ready for incoming requests and moving these instances between the inactive pools and an active state, ensuring that threading conditions within the bean are satisfied.
- Most importantly, automatically synchronizing data in an entity bean's instance variables with corresponding data items stored in persistent storage.

By dynamically maintaining a set of active bean instances and synchronizing bean state with persistent storage when beans are moved into and out of active state, the container makes it possible for an application to manage many more bean instances than could otherwise simultaneously be held in the application server's memory. In this respect, an EJB container provides services similar to virtual memory within an operating system.

Between transactions, the state of an entity bean can be cached. The EJB container supports option A, B, and C caching.

For more information about EJB containers, see "Resources for learning."

Managing EJB containers

Each application server can have a single EJB container; one is created automatically for you when the application server is created. The following steps are to be performed only as needed to improve performance after the EJB application has been deployed.

1. Adjust EJB container settings.
2. Adjust EJB cache settings.

If adjustments do not improve performance, consider adjusting access intent policies for entity beans, reassembling the module, and redeploying the module in the application.

EJB container settings

Use this page to configure and manage a specific EJB container.

To view this administrative console page, click **Servers > Application Servers > *serverName* > EJB Container**.

Passivation directory

Specifies the directory into which the container saves the persistent state of passivated stateful session beans.

Beans are passivated when the number of active bean instances becomes greater than the cache size specified in the container configuration. When a stateful bean is passivated, the container serializes the bean instance to a file in the passivation directory and discards the instance from the bean cache. If, at a later time, a request arrives for the passivated bean instance, the container retrieves it from the passivation directory, deserializes it, returns it to the cache, and dispatches the request to it. If any step fails (for example, if the bean instance is no longer in the passivation directory), the method invocation fails.

For a cluster of servers that span multiple systems in a sysplex and have stateful session beans with an activation policy of *Transaction* deployed in them, the passivation directory must reside on a hierarchical file system (HFS) that is shared across the multiple systems.

Inactive pool cleanup interval

Specifies the interval at which the container examines the pools of available bean instances to determine if some instances can be deleted to reduce memory usage.

Data type

Integer

Units	Milliseconds
Range	0 to 2 147 483 674

Default datasource JNDI name

Specifies the JNDI name of a data source to use if no data source is specified during application deployment. This setting is not applicable for EJB 2.x-compliant CMP beans.

Servlets and enterprise beans use *data sources* to obtain these connections. When configuring a container, you can specify a default data source for the container. This data source becomes the default data source used by any entity beans installed in the container that use container-managed persistence (CMP).

The default data source for a container is secure. When specifying it, you must provide a user ID and password for accessing the data source.

Specifying a default data source is optional if each CMP entity bean in the container has a data source specified in its configuration. If a default data source is not specified and a CMP entity bean is installed in the container without specifying a data source for that bean, applications cannot use that CMP entity bean.

Initial state

Specifies the execution state requested when the server first starts.

Data type	String
Default	Started
Range	Valid values are Started and Stopped

EJB container system properties

In addition to the settings accessible from the administrative console, you can set the following system property by command-line scripting:

com.ibm.websphere.ejbcontainer.poolSize

Specifies the size of the pool for the specified bean type. This property applies to stateless, message-driven and entity beans. If you do not specify a default value, the container defaults of 50 and 500 are used.

Set the pool size for a given entity bean as follows:

```
beantype=min,max[:beantype=min,max...]
```

beantype is the J2EE name of the bean, formed by concatenating the application name, the # character, the module name, the # character, and the name of the bean (that is, the string assigned to the <ejb-name> field in the bean's deployment descriptor). *min* and *max* are the minimum and maximum pool sizes, respectively, for that bean type. Do not specify the square brackets shown in the previous prototype; they denote optional additional bean types that you can specify after the first. Each bean-type specification is delimited by a colon (:).

Use an asterisk (*) as the value of *beantype* to indicate that all bean types are to use those values unless overridden by an exact bean-type specification somewhere else in the string, as follows:

```
*=30,100
```


To specify that a default value be used, omit either *min* or *max* but retain the comma (,) between the two values, as follows (split for publication):

```
SMAApp#PerfModule#TunerBean=54,  
:SMAApp#SMModule#TypeBean=100,200
```

You can specify the bean types in any order within the string.

EJB cache settings

Use this page to configure and manage the cache for a specific EJB container. To determine the cache absolute limit, multiply the number of enterprise beans active in any given transaction by the total number of concurrent transactions expected. Then, add the number of active session bean instances. You can use the Tivoli Performance Viewer to view bean performance information.

To view this administrative console page, click **Servers > Application Servers > *serverName* > EJB Container > EJB Cache Settings**.

Cleanup interval

Specifies the interval at which the container attempts to remove unused items from the cache in order to reduce the total number of items to the value of the cache size.

The cache manager tries to maintain some unallocated entries that can be allocated quickly as needed. A background thread attempts to free some entries while maintaining some unallocated entries. If the thread runs while the application server is idle, when the application server needs to allocate new cache entries, it does not pay the performance cost of removing entries from the cache. In general, increase this parameter as the cache size increases.

Data type	Integer
Units	Milliseconds
Range	0 to 2 147 483 674
Default	3000

Cache size

Specifies the number of buckets in the active instance list within the EJB container.

A bucket can contain more than one active enterprise bean instance, but performance is maximized if each bucket in the table has a minimum number of instances assigned to it. When the number of active instances within the container exceeds the number of buckets, that is, the cache size, the container periodically attempts to reduce the number of active instances in the table by passivating some of the active instances. For the best balance of performance and memory, set this value to the maximum number of active instances expected during a typical workload.

Data type	Integer
Units	Buckets in the hash table
Range	Greater than 0. The container selects the next largest prime number equal to or greater than the specified value.
Default	2053

Container interoperability

Container interoperability describes the ability of WebSphere Application Server clients and servers at different versions to successfully negotiate differences in native Enterprise JavaBeans (EJB) Version 1.1 finder methods support and Java 2 Platform, Enterprise Edition (J2EE) Version 1.3 compliance.

At one time, there were significant interoperability problems among WebSphere Application Server, versions 4.0.x and 3.5.x distributed, and Version 4.0.x for zSeries. The introduction of interoperable versions of some class types solved these problems for distributed versions 3.5.6, 4.0.3, and 5 as well as for zSeries Version 4.0.x.

Older 4.0.x and 3.5.x client and application server versions do not support the interoperability classes, which makes them uninteroperable with versions that use the classes. The system property `com.ibm.websphere.container.portable` remedies this situation by enabling newer versions of the application server to turn off the interoperability classes. This lets a more recent application server return class types that are interoperable with an older client.

Depending on the value of `com.ibm.websphere.container.portable`, application servers at versions 5 and later, 4.0.3 and later, and 3.5.6 and later, return different classes for the following:

- Enumerations and collections returned by EJB 1.1 finder methods
- EJBMetaData
- Handles to:
 - Entity beans
 - Session beans
 - Home interfaces

If the property is set to `false`, application servers return the old class types, to enable interoperability with versions 3.5.5 and earlier, and 4.0.2 and earlier. If the property is set to `true`, application servers return the new classes.

Instructions for setting the `com.ibm.websphere.container.portable` property are in the release notes for versions 3.5.6 and later, and 4.0.3 and later. The following tables show interoperability characteristics for various version combinations of application servers and clients as well as default property values for each combination.

Interoperability of Version 3.5.x client with Version 5 (and later) application server

Clients at Version 3.5.5 and earlier are not interoperable with Version 5 and later servers when using:

- EJBMetaData
- Enumerations returned by EJB 1.x finder methods
- Handles to entity beans

If you would like to use updated Handle classes in EJB 2.x-compliant beans but have one of the older clients (versions 3.5.5 and earlier) installed, set the system property `com.ibm.websphere.container.portable.finder` to `false`. With this setting in place, the Version 5 application server uses the updated handles but returns the enumerations and collections that were used in the earlier clients.

To interoperate with Version 5 application servers, you must upgrade all Version 3.5.x clients to Version 3.5.6 or later.

Interoperability of Version 5 (and later) client with Version 3.5.x application server

Client at Version 5 and later, using this function	Application server at Version 3.5.6, property true	Application server at Version 3.5.6, property false (default)	Application server at Version 3.5.5 and earlier
EJBMetaData	Does not work across domains	Works	Does not work
Handle to session bean	Works	Works	Does not work
Handle to entity bean	Does not work across domains	Does not work across domains	Does not work across domains
Enumeration returned by EJB 1.x finder method	Works	Works	Works

Interoperability of Version 4.0.x client with Version 5 (and later) application server

Ideally, all 4.0.x clients that use Version 5 application servers should be at Version 4.0.3 or later.

Version 5 and later application servers return the interoperability class types by default (true). This can cause interoperability problems for distributed clients at versions 4.0.1 or 4.0.2. In particular, problems can occur with collections and enumerations returned by EJB 1.1 finder methods.

Although it is strongly discouraged, you can set `com.ibm.websphere.container.portable` to `false` on a Version 5 and later application server. This causes the application server to return the old class types, providing interoperability with clients at Version 4.0.2 and earlier. This is discouraged because:

- The Version 5 application server instance would become non-J2EE 1.3 compliant with regard to handles, home interface handles, and EJBMetaData.
- EJB 1.x finder methods return collection and enumeration objects that do not originate from `ejbportable.jar`.
- Interoperability restrictions still exist with the property set to `false`.
- Version 5 and later client handles to entity beans and home interfaces do not work across domains for the server you set to `false`.

If you would like to use updated Handle classes in EJB 2.x-compliant beans but have one of the older clients (versions 4.0.2 and earlier) installed, set the system property `com.ibm.websphere.container.portable.finder` to `false`. With this setting in place, the Version 5 and later application server uses the updated handles but returns the enumerations and collections that were used in the earlier clients.

Interoperability of client at Version 4.0.2 and earlier with Version 5 (and later) application server

Client at Version 4.0.2 and earlier, using this function	Application server at Version 5 and later, property true (default)	Application server at Version 5 and later, property false
EJBMetaData	Does not work	Works for 4.0.2 client
Handle to session bean	Does not work	Works

Client at Version 4.0.2 and earlier, using this function	Application server at Version 5 and later, property true (default)	Application server at Version 5 and later, property false
Handle to entity bean	Does not work	Does not work across cells
Enumeration returned by EJB 1.x finder method	Does not work	Works
Collection returned by EJB 1.x finder method	Does not work	Works
Handle to home interface	Does not work	Does not work across cells

If you would like to use updated Handle classes in EJB 2.x-compliant beans but have one of the older clients (versions 3.5.5 and earlier, and 4.0.2 and earlier) installed, set the system property `com.ibm.websphere.container.portable.finder>false`. With this setting in place, the Version 5 and later server uses the new Handle classes but returns the older enumeration and collection classes.

Interoperability of client at Version 4.0.3 and later with Version 5 and later application server

Clients at Version 4.0.3 and later work well with Version 5 and later application servers. However, if you set the `com.ibm.websphere.container.portable` to `false`, client handles to entity beans and home interfaces do not work across domains for the server you set to `false`.

Client at Version 4.0.3 and later, using this function	Application server at Version 5 and later, property true (default)	Application server at Version 5 and later, property false
EJBMetaData	Works	Works
Handle to session bean	Works	Works
Handle to entity bean	Works	Does not work across cells
Enumeration returned by EJB 1.x finder method	Works	Works
Collection returned by EJB 1.x finder method	Works	Works
Handle to home interface	Works	Does not work across cells

Interoperability of Version 5 and later client with Version 4.0.x application server

Clients at Version 5 and later work well with Version 4.0.3 application servers if you set `com.ibm.websphere.container.portable` to `true`. Client handles to entity beans and home interfaces do not work across domains for any Version 4.0.3 server with `com.ibm.websphere.container.portable` at the default value, `false`. Version 5 client handles to application servers at Version 4.0.2 and earlier also have restrictions.

Client at Version 5 and later, using this function	Application server at Version 4.0.3, property true	Application server at Version 4.0.3, property false (default)	Application server at Version 4.0.2 or earlier
EJBMetaData	Works	Works	Works for 4.0.2 server only

Client at Version 5 and later, using this function	Application server at Version 4.0.3, property true	Application server at Version 4.0.3, property false (default)	Application server at Version 4.0.2 or earlier
Handle to session bean	Works	Works	Works
Handle to entity bean	Works	Does not work across domains	Does not work across domains
Enumeration returned by EJB 1.x finder method	Works	Works	Works
Collection returned by EJB 1.x finder method	Works	Works	Works
Handle to home interface	Works	Does not work across domains	Does not work across domains

Interoperability of zSeries Version 4.0.x client with Version 5 and later application server

The only valid configuration for container interoperability with zSeries Version 4.0.x clients is the default configuration for the Version 5 application server.

Interoperability of Version 5 and later client with zSeries Version 4.0.x application server

Version 5 clients should work with a zSeries Version 4.0.x application server with the correct interoperability fixes described in the zSeries documentation. The interoperability characteristics should be the same as for a Version 4.0.3 distributed application server with the property set to true.

Client at Version 5 and later, using this function	zSeries application server at Version 4.0.x
EJBMetaData	Works
Handle to session bean	Works
Handle to entity bean	Works
Enumeration returned by EJB 1.x finder method	Works
Collection returned by EJB 1.x finder method	Works
Handle to home interface	Works

Interoperability of the handle formats in WebSphere Application Server, Version 5 and Version 5.0.1

Applications that attempt to persist handles to enterprise beans and **EJBHome** needed to subclass **ObjectInputStream** in WebSphere Application Server, Version 5. This action was required so that the subclass **ObjectInputStream** could utilize the context class loader to resolve the classes for enterprise beans and **EJBHome** stubs.

In addition, handles created and persisted in WebSphere Application Server, Version 5 only work with objects that have an unchanged remote interface. If the

remote interface is changed, the handle is no longer valid because the stub is serialized inside the handle and its serial Version UID changes if the remote interface changes.

This release introduces a new handle persistence mechanism that avoids the implementation drawbacks of the previous version. However, if handles are used for this WebSphere Application Server deployment, you should consider the following issues when applying this update, future WebSphere Application Server Fix Packs and EJB Container cumulative fixes for WebSphere Application Server, Version 5.

If a WebSphere Application Server, Version 5 persisted handle or home handle is encountered by a WebSphere Application Server, Version 5.0.1 system, it can be read and utilized. In addition, it will be converted to WebSphere Application Server, Version 5.0.1 format if it is re-persisted. The WebSphere Application Server, Version 5.0.1 format cannot be read by a WebSphere Application Server, Version 5 system unless PQ72184 is applied.

Problems arise when handles are persisted and shared across systems that are not at the WebSphere Application Server, Version 5.0.1 level or later. However, a Version 5 system can receive a handle from Version 5.0.1 remotely through a call to get a handle on an enterprise bean or a getHomeHandle on an **EJBHome**. The remote call will succeed, however, any attempt to persist it on the Version 5 system will have the same limitations regarding the use of `ObjectInputStream` and changes in remote interface invalidating the persisted handle.

When your application stores handles persistently and shares this persistence with multiple clients or application servers, apply WebSphere Application Server, Version 5.0.1 or PQ72184 to both the client and server systems at the same time. Failure to do so can result in the inability of these systems to read the handle data stored by upgraded systems. Also, handles stored by the WebSphere Application Server, Version 5 can force the applications of the updated system to still subclass `ObjectInputStream`. Applications using the WebSphere Application Server Enterprise, Version 5 scheduler and process choreographer, are affected by these changes. These users should update their Version 5 systems at the same time with either Version 5.0.1 or PQ72184.

If the applications store handles in the session context, or locally in a file on the same system, that is not shared by other applications, on different systems, they might be able to update their systems individually, rather than all at once. If Client Container and thin client applications do not share persisted handle data, they can be updated as needed as well. However, handles created and persisted in WebSphere Application Server, Version 5, Version 4.0.3 and later (with the property flag set), or Version 3.5.7 and later (with the property flag set) are not usable if either the home or the remote interface changes.

If any WebSphere Application Server, Version 3.5.7 or Version 4.0.3 and later enables the system property `com.ibm.websphere.container.portable` to **true**, any handles to objects on that server have the same interoperability limitations. In addition, if any WebSphere Application Server, Version 3.5.7 and later or Version 4.0.3 applications store a handle obtained from a WebSphere Application Server, Version 5 or Version 5.0.1, the same restrictions apply, regarding the need to subclass `ObjectInputStream` and the usability of handles after a change to the remote interface is made.

Replication of the Http Session and Handles

This note applies to you if you place Handles to Homes or EJBs, or EJB or EJBHome references in the Http Session in your application and you use Http Session Replication. If you intend to replicate a mixed environment of Version 5.0.0 and Version 5.0.1 or 5.0.2 machines you should first apply the latest Version 5.0.0 container cumulative e-fix to the Version 5.0.0 machines before allowing the Version 5.0.1 or 5.0.2 server into the typology. The reason for this is that Version 5.0.0 servers are not able to understand the persisted Handle format used on the Version 5.0.1 and 5.0.2 server. This is similar to the case of Version 5.0.0 and Version 5.0.1 or 5.0.2 systems trying to use a shared database, mentioned above. But in this case, it is the Http Session object and not the database providing the persistence.

Top Down Deployment Mapping

The size of the Handle objects has grown due to the fix put in to allow serialization and deserialization to occur without the previous requirements of subclassing the ObjectInputStream and so on. Top down deployment of an object that contains EJB and EJBHome references create a database table ddl that has a field of 1000 bytes of VARCHAR for BITDATA which will contain the Handle. It might be that your object's Handle does not fit in the 1000 byte default field, and you might need to adjust this to a higher value. You might try increments of 250 bytes, that is, 1250, 1500, and so on.

Deploying EJB modules

Assemble one or more EJB modules, assemble one or more Web modules, and assemble them into a J2EE application.

1. Prepare the deployment environment.
2. Deploy the application.
3. **5.1+** Update the configuration for each EJB module as needed for the deployment environment.
4. For information about the EJB deployment tool, see the EJB deployment tool.

The next step is to test and debug the module.

EJB module collection

Use this page to manage the EJB modules deployed in a specific application.

To view this administrative console page, click **Applications > Enterprise Applications > *applicationName* > EJB modules**. Click the check boxes to select one or more of the EJB modules in your collection.

URI

When resolved relative to the application URL, this specifies the location of the module's archive contents on a file system. The URI matches the <ejb> or <web> tag in the <module> tag of the application deployment descriptor.

EJB module settings

Use this page to configure and manage a specific deployed EJB module.

Note: You cannot start or stop an individual EJB module for modification. You must start or stop the appropriate application entirely.

To view this administrative console page, click **Applications > Enterprise Applications > *applicationName* > EJB modules > *moduleName***.

URI

When resolved relative to the application URL, this specifies the location of the module archive contents on a file system. The URI must match the URI of a ModuleRef URI in the deployment descriptor of the deployed application (EAR).

Alternate DD

Specifies a deployment descriptor to be used at run time instead of the one installed in the module.

Starting weight

Specifies the order in which modules are started when the server starts. The module with the lowest starting weight is started first.

Data type	Integer
Default	5000
Range	Greater than 0

Enterprise beans: Resources for learning

Use the following links to find relevant supplemental information about enterprise beans. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to this product but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Planning, business scenarios, and IT architecture
- Programming model and decisions
- Programming instructions and examples
- Programming specifications

Planning, business scenarios, and IT architecture

- Mastering Enterprise JavaBeans

A comprehensive treatment of Enterprise JavaBeans (EJB) programming in nonprintable form (PDF). One must be registered to download the PDF, but registration is free. Information about purchasing a hardcopy is available on the Web site.

- *Enterprise JavaBeans* by Richard Monson-Haefel (O'Reilly and Associates, Inc.: Third Edition, 2001)

Programming model and decisions

- Read all about EJB 2.0

A comprehensive overview of the specification.

- The J2EE Tutorial

This set of articles by Sun Microsystems covers several EJB-related topics, including the basic programming models, persistence, and EJB Query Language.

Programming instructions and examples

- Rules and Patterns for Session Facades
EJB programming practice: Fronting entity beans with a session-bean facade.
- WebSphere Application Server Development Best Practices for Performance and Scalability
Programming practice for enterprise beans and other types of J2EE components.
- Optimistic Locking in IBM WebSphere Application Server 4.0.2
Examples of the effect of optimistic concurrency on application behavior.
Although the paper is based on a previous version of this product, the data access issues discussed in it are current.
This paper does not seem to be available directly by URL. To view this paper, visit the specified URL and search on "optimistic locking"

Programming specifications

- What's new in the Enterprise JavaBeans 2.0 Specification?
You can also download the specification itself from this URL.
- Java™ 2 Platform: Compatibility with Previous Releases
This Sun Microsystems article includes both source and binary compatibility issues.

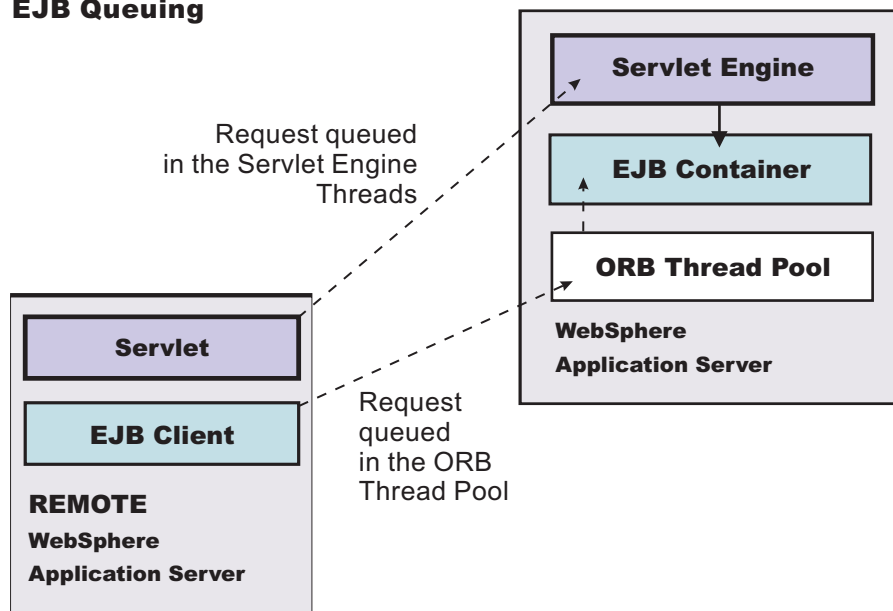
EJB method Invocation Queuing

Method invocations to enterprise beans are only queued for remote clients, making the method call. An example of a remote client is an enterprise Java bean (EJB) client running in a separate Java virtual machine (JVM) (another address space) from the enterprise bean. In contrast, no queuing occurs if the EJB client, either a servlet or another enterprise bean, is installed in the same JVM on which the EJB method runs and on the same thread of execution as the EJB client.

Remote enterprise beans communicate by using the Remote Method Invocation over an Internet Inter-Orb Protocol (RMI-IIOP). Method invocations initiated over RMI-IIOP are processed by a server-side object request broker (ORB). The thread pool acts as a queue for incoming requests. However, if a remote method request is issued and there are no more available threads in the thread pool, a new thread is created. After the method request completes the thread is destroyed. Therefore, when the ORB is used to process remote method requests, the EJB container is an open queue, due to the use of unbounded threads. The following illustration

depicts the two queuing options of enterprise beans.

EJB Queuing




The following are two tips for queueing enterprise beans:


- **Analyze the calling patterns of the EJB client.**

When configuring the thread pool, it is important to understand the calling patterns of the EJB client. If a servlet is making a small number of calls to remote enterprise beans and each method call is relatively quick, consider setting the number of threads in the ORB thread pool to a value lower than the Web container thread pool size value.

Short-lived EJB calls



Longer-lived EJB calls



The degree to which the ORB thread pool value needs increasing is a function of the number of simultaneous servlets, that is, clients, calling enterprise beans and the duration of each method call. If the method calls are longer or the applications spend a lot of time in the ORB, consider making the ORB thread pool size equal to the Web container size. If the servlet makes only short-lived or quick calls to the ORB, servlets can potentially reuse the same ORB thread. In this case, the ORB thread pool can be small, perhaps even one-half of the thread pool size setting of the Web container.

- **Monitor the percentage of configured threads in use.**

Tivoli Performance Viewer shows a metric called *percent maxed*, which is used to determine how often the configured threads are used. A value that is consistently in the double-digits, indicates a possible bottleneck at the ORB. Increase the number of threads.

See also [Queuing network](#).

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

Trademarks and service marks

The following terms are trademarks of IBM Corporation in the United States, other countries, or both:

- AIX
- AS/400
- CICS
- Cloudscape
- DB2
- DFSMS
- Domino
- Everyplace
- iSeries
- IBM
- IMS
- Informix
- iSeries
- Language Environment
- Lotus
- MQSeries
- MVS
- OS/390
- RACF
- Redbooks
- RMF
- SecureWay
- SupportPac
- Tivoli
- ViaVoice
- VisualAge
- VTAM
- WebSphere
- z/OS
- zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, 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.