



# BEA WebLogic Server™

## Programming WebLogic JMS

## Copyright

Copyright © 2002 BEA Systems, Inc. All Rights Reserved.

## Restricted Rights Legend

This software and documentation is subject to and made available only pursuant to the terms of the BEA Systems License Agreement and may be used or copied only in accordance with the terms of that agreement. It is against the law to copy the software except as specifically allowed in the agreement. This document may not, in whole or in part, be copied photocopied, reproduced, translated, or reduced to any electronic medium or machine readable form without prior consent, in writing, from BEA Systems, Inc.

Use, duplication or disclosure by the U.S. Government is subject to restrictions set forth in the BEA Systems License Agreement and in subparagraph (c)(1) of the Commercial Computer Software-Restricted Rights Clause at FAR 52.227-19; subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013, subparagraph (d) of the Commercial Computer Software--Licensing clause at NASA FAR supplement 16-52.227-86; or their equivalent.

Information in this document is subject to change without notice and does not represent a commitment on the part of BEA Systems. THE SOFTWARE AND DOCUMENTATION ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. FURTHER, BEA Systems DOES NOT WARRANT, GUARANTEE, OR MAKE ANY REPRESENTATIONS REGARDING THE USE, OR THE RESULTS OF THE USE, OF THE SOFTWARE OR WRITTEN MATERIAL IN TERMS OF CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE.

## Trademarks or Service Marks

BEA, Jolt, Tuxedo, and WebLogic are registered trademarks of BEA Systems, Inc. BEA Builder, BEA Campaign Manager for WebLogic, BEA eLink, BEA Manager, BEA WebLogic Commerce Server, BEA WebLogic Enterprise, BEA WebLogic Enterprise Platform, BEA WebLogic Express, BEA WebLogic Integration, BEA WebLogic Personalization Server, BEA WebLogic Platform, BEA WebLogic Portal, BEA WebLogic Server, BEA WebLogic Workshop and How Business Becomes E-Business are trademarks of BEA Systems, Inc.

All other trademarks are the property of their respective companies.

Programming WebLogic JMS

<b>Part Number</b>	<b>Document Revised</b>	<b>Software Version</b>
N/A	August 23, 2002	BEA WebLogic Server Version 7.0

---

# Contents

## About This Document

Audience.....	xii
e-docs Web Site.....	xii
How to Print the Document.....	xii
Related Information.....	xiii
Contact Us!.....	xiii
Documentation Conventions.....	xiv

## 1. Introduction to WebLogic JMS

What Is JMS?.....	1-1
Implementation of Java Specifications.....	1-2
J2EE Specification.....	1-2
JMS Specification.....	1-2
WebLogic JMS Features.....	1-3
WebLogic JMS Architecture.....	1-3
Major Components.....	1-4
Clustering Features.....	1-5
WebLogic JMS Extensions.....	1-7
JMS Enhancements in WebLogic Server 7.0.....	1-8
High Availability Enhancement.....	1-8
Distributed Destinations within a WebLogic Cluster.....	1-8
Flow Control.....	1-9
WebLogic Messaging Bridge.....	1-9
Message Paging.....	1-9

## 2. WebLogic JMS Fundamentals

Messaging Models.....	2-2
-----------------------	-----

---

Point-to-Point Messaging .....	2-2
Publish/Subscribe Messaging.....	2-3
Message Persistence .....	2-4
WebLogic JMS Classes .....	2-5
ConnectionFactory.....	2-6
Connection.....	2-8
Session .....	2-9
Non-transacted Session .....	2-10
Transacted Session .....	2-12
Destination.....	2-13
Distributed Destination.....	2-14
MessageProducer and MessageConsumer.....	2-14
Message .....	2-16
Message Header Fields.....	2-16
Message Property Fields.....	2-20
Message Body .....	2-21
ServerSessionPoolFactory .....	2-22
ServerSessionPool .....	2-22
ServerSession.....	2-23
ConnectionConsumer .....	2-23

### **3. Managing WebLogic JMS**

Configuring WebLogic JMS .....	3-1
Configuring WebLogic JMS Clustering.....	3-3
How JMS Clustering Works.....	3-3
JMS Clustering Naming Requirements.....	3-4
JMS Distributed Destination within a Cluster.....	3-4
JMS as a Migratable Service within a Cluster .....	3-5
Configuration Guidelines for JMS Clustering.....	3-5
What About Failover? .....	3-6
Configuring JMS Migratable Targets.....	3-6
How JMS Migration Works .....	3-7
Configuration Steps for JMS Server Migration.....	3-7
Persistent Store Migration.....	3-8
Migration Failover.....	3-9

Tuning WebLogic JMS .....	3-9
Monitoring WebLogic JMS.....	3-10
Recovering from a WebLogic Server Failure .....	3-11

## 4. Developing a WebLogic JMS Application

Application Development Flow .....	4-2
Importing Required Packages.....	4-3
Setting Up a JMS Application.....	4-4
Step 1: Look Up a Connection Factory in JNDI .....	4-6
Step 2: Create a Connection Using the Connection Factory .....	4-7
Create a Queue Connection.....	4-7
Create a Topic Connection.....	4-7
Step 3: Create a Session Using the Connection .....	4-8
Create a Queue Session.....	4-8
Create a Topic Session.....	4-9
Step 4: Look Up a Destination (Queue or Topic) .....	4-9
Server Affinity When Looking Up Destinations .....	4-10
Step 5: Create Message Producers and Message Consumers Using the Session and Destinations.....	4-11
Create QueueSenders and QueueReceivers .....	4-11
Create TopicPublishers and TopicSubscribers .....	4-12
Step 6a: Create the Message Object (Message Producers) .....	4-13
Step 6b: Optionally Register an Asynchronous Message Listener (Message Consumers) .....	4-15
Step 7: Start the Connection.....	4-16
Example: Setting Up a PTP Application.....	4-16
Example: Setting Up a Pub/Sub Application .....	4-19
Sending Messages .....	4-22
Step 1: Create a Message Object.....	4-22
Step 2: Define a Message .....	4-22
Step 3: Send the Message to a Destination .....	4-23
Send a Message Using Queue Sender.....	4-23
Send a Message Using TopicPublisher.....	4-25
Setting Message Producer Attributes .....	4-26
Example: Sending Messages Within a PTP Application .....	4-28

---

Example: Sending Messages Within a Pub/Sub Application.....	4-28
Receiving Messages .....	4-29
Receiving Messages Asynchronously .....	4-30
Asynchronous Message Pipeline.....	4-30
Receiving Messages Synchronously .....	4-31
Example: Receiving Messages Synchronously Within a PTP Application	
4-32	
Example: Receiving Messages Synchronously Within a Pub/Sub	
Application.....	4-32
Recovering Received Messages .....	4-33
Acknowledging Received Messages .....	4-33
Releasing Object Resources .....	4-34
Managing Rolled Back, Recovered, or Expired Messages.....	4-35
Setting a Redelivery Delay for Messages.....	4-35
Setting a Redelivery Delay.....	4-36
Overriding the Redelivery Delay on a Destination .....	4-36
Setting a Redelivery Limit for Messages .....	4-37
Configuring a Message Redelivery Limit .....	4-37
Configuring an Error Destination for Undelivered Messages.....	4-37
Passive Message Expiration Policy .....	4-38
Setting Message Delivery Times .....	4-38
Setting a Delivery Time on Producers.....	4-39
Setting a Delivery Time on Messages .....	4-39
Overriding a Delivery Time .....	4-40
Interaction with the Time-to-Live Value .....	4-40
Setting a Relative Time-to-Deliver Override .....	4-40
Setting a Scheduled Time-to-Deliver Override.....	4-41
JMS Schedule Interface.....	4-43
Managing Connections .....	4-44
Defining a Connection Exception Listener .....	4-44
Accessing Connection Metadata .....	4-45
Starting, Stopping, and Closing a Connection.....	4-46
Managing Sessions .....	4-47
Defining a Session Exception Listener.....	4-48
Closing a Session.....	4-49

---

Creating Destinations Dynamically.....	4-49
Using the JMSHelper Class Methods.....	4-50
Using Temporary Destinations.....	4-52
Creating a Temporary Queue.....	4-52
Creating a Temporary Topic.....	4-53
Deleting a Temporary Destination.....	4-53
Setting Up Durable Subscriptions.....	4-53
Defining the Persistent Store.....	4-54
Defining the Client ID.....	4-55
Creating Subscribers for a Durable Subscription.....	4-56
Deleting Durable Subscriptions.....	4-57
Modifying Durable Subscriptions.....	4-57
Managing Durable Subscriptions.....	4-58
Setting and Browsing Message Header and Property Fields.....	4-58
Setting Message Header Fields.....	4-59
Setting Message Property Fields.....	4-61
Browsing Header and Property Fields.....	4-65
Filtering Messages.....	4-66
Defining Message Selectors Using SQL Statements.....	4-67
Defining XML Message Selectors Using XML Selector Method.....	4-68
Displaying Message Selectors.....	4-69
Indexing Topic Subscriber Message Selectors To Optimize Performance.....	4-69
Defining Server Session Pools.....	4-71
Step 1: Look Up Server Session Pool Factory in JNDI.....	4-73
Step 2: Create a Server Session Pool Using the Server Session Pool Factory.	4-74
Create a Server Session Pool for Queue Connection Consumers.....	4-74
Create a Server Session Pool for Topic Connection Consumers.....	4-75
Step 3: Create a Connection Consumer.....	4-75
Create a Connection Consumer for Queues.....	4-76
Create a Connection Consumer for Topics.....	4-76
Example: Setting Up a PTP Client Server Session Pool.....	4-77
Example: Setting Up a Pub/Sub Client Server Session Pool.....	4-79
Using Multicasting.....	4-81

---

Step 1: Set Up the JMS Application, Creating Multicast Session and Topic Subscriber.....	4-83
Step 2: Set Up the Message Listener.....	4-84
Dynamically Configuring Multicasting Configuration Attributes .....	4-85
Example: Multicast TTL .....	4-86
Using Distributed Destinations.....	4-88
Accessing Distributed Destinations.....	4-88
Looking Up Distributed Queues .....	4-89
Looking Up Distributed Topics.....	4-91
Accessing Distributed Destination Members .....	4-94
Load Balancing Messages Across a Distributed Destination.....	4-94
Load Balancing Options.....	4-94
Consumer Load Balancing.....	4-96
Producer Load Balancing .....	4-96
Load Balancing Heuristics .....	4-96
Defeating Load Balancing.....	4-98
How Distributed Destination Load Balancing Is Affected When Using the “Server Affinity Enabled” Attribute .....	4-99
Distributed Destination Migration.....	4-101
Distributed Destination Failover .....	4-102

## **5. Using Transactions with WebLogic JMS**

Overview of Transactions.....	5-1
Using JMS Transacted Sessions .....	5-3
Step 1: Set Up JMS Application, Creating Transacted Session .....	5-4
Step 2: Perform Desired Operations.....	5-4
Step 3: Commit or Roll Back the JMS Transacted Session .....	5-4
Using JTA User Transactions.....	5-5
Step 1: Set Up JMS Application, Creating Non-Transacted Session .....	5-7
Step 2: Look Up User Transaction in JNDI .....	5-7
Step 3: Start the JTA User Transaction .....	5-7
Step 4: Perform Desired Operations.....	5-8
Step 5: Commit or Roll Back the JTA User Transaction .....	5-8
Asynchronous Messaging Within JTA User Transactions Using Message Driven Beans .....	5-9
Example: JMS and EJB in a JTA User Transaction.....	5-9

---

## 6. Porting WebLogic JMS Applications

Existing Feature Functionality Changes.....	6-1
Existing Feature 5.1 to 6.0 Functionality Changes .....	6-2
Existing Feature 6.0 to 6.1 Functionality Changes .....	6-7
Porting Existing Applications.....	6-8
Before You Begin.....	6-8
Porting Steps for 4.5 and 5.1 Applications to 6.x .....	6-9
Porting Steps for 6.0 Applications to 6.1 .....	6-11
Porting Steps for 6.x Applications to 7.0 .....	6-11
Deleting JDBC Database Stores.....	6-12

### A. Configuration Checklists

Server Clusters.....	A-2
JTA User Transactions .....	A-2
JMS Transactions .....	A-2
Message Delivery .....	A-3
Asynchronous Message Delivery .....	A-3
Persistent Messages .....	A-4
Concurrent Message Processing.....	A-4
Multicasting.....	A-5
Durable Subscriptions .....	A-5
Destination Sort Order.....	A-6
Temporary Destinations .....	A-6
Thresholds and Quotas .....	A-7

### B. JDBC Database Utility

Overview .....	B-1
About JMS Tables .....	B-1
Regenerating JDBC Database Stores .....	B-2



---

# About This Document

This document explains how to use the BEA WebLogic Server™ platform to implement the Java™ Messaging Service (JMS) API for accessing enterprise messaging systems.

The document is organized as follows:

- [Chapter 1, “Introduction to WebLogic JMS,”](#) provides an overview of WebLogic Java Message Service (JMS).
- [Chapter 2, “WebLogic JMS Fundamentals,”](#) describes WebLogic JMS components and features.
- [Chapter 3, “Managing WebLogic JMS,”](#) provides an overview of configuring and monitoring WebLogic JMS.
- [Chapter 4, “Developing a WebLogic JMS Application,”](#) describes how to develop a WebLogic JMS application.
- [Chapter 5, “Using Transactions with WebLogic JMS,”](#) describes how to use transactions with WebLogic JMS.
- [Chapter 6, “Porting WebLogic JMS Applications,”](#) describes how to port your WebLogic JMS applications to a new release of WebLogic Server.
- [Appendix A, “Configuration Checklists,”](#) provides monitoring checklists for various WebLogic JMS features.
- [Appendix B, “JDBC Database Utility,”](#) describes how to use the JDBC database utility to generate new JDBC stores and delete existing ones.

---

# Audience

This document is written for application developers who want to design, develop, configure, and manage JMS applications using the Java 2 Platform, Enterprise Edition (J2EE) from Sun Microsystems. It is assumed that readers know JMS, JNDI (Java Naming and Directory Interface), the Java programming language, the Enterprise JavaBeans™ (EJB™), and the Java Transaction API (JTA) of the J2EE specification.

## e-docs Web Site

BEA product documentation is available on the BEA corporate Web site. From the BEA Home page, click on Product Documentation. Or you can go directly to the WebLogic Server Product Documentation page at <http://edocs.bea.com/wls/docs70>.

## How to Print the Document

You can print a copy of this document from a Web browser, one main topic at a time, by using the File→Print option on your Web browser.

A PDF version of this document is available on the WebLogic Server documentation Home page on the e-docs Web site (and also on the documentation CD). You can open the PDF in Adobe Acrobat Reader and print the entire document (or a portion of it) in book format. To access the PDFs, open the WebLogic Server documentation Home page, click Download Documentation, and select the document you want to print.

Adobe Acrobat Reader is available at no charge from the Adobe Web site at <http://www.adobe.com>.

---

# Related Information

The BEA corporate Web site provides all documentation for WebLogic Server. For more information on JMS, access the JMS Specification and Javadoc supplied on Sun Microsystems' Java Web site at the follow location:

<http://http://java.sun.com/products/jms/docs.html>

## Contact Us!

Your feedback on BEA documentation is important to us. Send us e-mail at [docsupport@bea.com](mailto:docsupport@bea.com) if you have questions or comments. Your comments will be reviewed directly by the BEA professionals who create and update the documentation.

In your e-mail message, please indicate the software name and version you are using, as well as the title and document date of your documentation. If you have any questions about this version of BEA WebLogic Server, or if you have problems installing and running BEA WebLogic Server, contact BEA Customer Support through BEA WebSupport at <http://www.bea.com>. You can also contact Customer Support by using the contact information provided on the Customer Support Card, which is included in the product package.

When contacting Customer Support, be prepared to provide the following information:

- Your name, e-mail address, phone number, and fax number
- Your company name and company address
- Your machine type and authorization codes
- The name and version of the product you are using
- A description of the problem and the content of pertinent error messages

---

# Documentation Conventions

The following documentation conventions are used throughout this document.

<b>Convention</b>	<b>Usage</b>
Ctrl+Tab	Keys you press simultaneously.
<i>italics</i>	Emphasis and book titles.
monospace text	Code samples, commands and their options, Java classes, data types, directories, and file names and their extensions. Monospace text also indicates text that you enter from the keyboard. <i>Examples:</i> <pre>import java.util.Enumeration; chmod u+w * config/examples/applications .java config.xml float</pre>
<i>monospace italic text</i>	Variables in code. <i>Example:</i> <pre>String CustomerName;</pre>
UPPERCASE TEXT	Device names, environment variables, and logical operators. <i>Examples:</i> <pre>LPT1 BEA_HOME OR</pre>
{ }	A set of choices in a syntax line.
[ ]	Optional items in a syntax line. <i>Example:</i> <pre>java utils.MulticastTest -n name -a address [-p portnumber] [-t timeout] [-s send]</pre>

---

<b>Convention</b>	<b>Usage</b>
	Separates mutually exclusive choices in a syntax line. <i>Example:</i> <pre>java weblogic.deploy [list deploy undeploy update] password {application} {source}</pre>
...	Indicates one of the following in a command line: <ul style="list-style-type: none"><li>■ An argument can be repeated several times in the command line.</li><li>■ The statement omits additional optional arguments.</li><li>■ You can enter additional parameters, values, or other information</li></ul>
.	Indicates the omission of items from a code example or from a syntax line.

---



# 1 Introduction to WebLogic JMS

The following sections provide an overview of the Java Message Service (JMS) for BEA WebLogic Server:

- [What Is JMS?](#)
- [Implementation of Java Specifications](#)
- [WebLogic JMS Features](#)
- [WebLogic JMS Architecture](#)
- [WebLogic JMS Extensions](#)
- [JMS Enhancements in WebLogic Server 7.0](#)

## What Is JMS?

An enterprise messaging system, also referred to as Message-Oriented Middleware (MOM), enables applications to communicate with one another through the exchange of messages. A message is a request, report, and/or event that contains information needed to coordinate communication between different applications. A message provides a level of abstraction, allowing you to separate the details about the destination system from the application code.

The Java Message Service (JMS) is a standard API for accessing enterprise messaging systems. Specifically, JMS:

- Enables Java applications sharing a messaging system to exchange messages.
- Simplifies application development by providing a standard interface for creating, sending, and receiving messages.

The following figure illustrates WebLogic JMS messaging.

**Figure 1-1 WebLogic JMS Messaging**



As illustrated in the figure, WebLogic JMS accepts messages from *producer* applications and delivers them to *consumer* applications.

## Implementation of Java Specifications

WebLogic Server is compliant with the following Java specifications.

### J2EE Specification

WebLogic Server 7.0 is compliant with Sun Microsystems' J2EE 1.3 specification.

### JMS Specification

WebLogic Server 7.0 is fully compliant with the [JMS Specification – version 1.0.2b](#) and can be used in production.

# WebLogic JMS Features

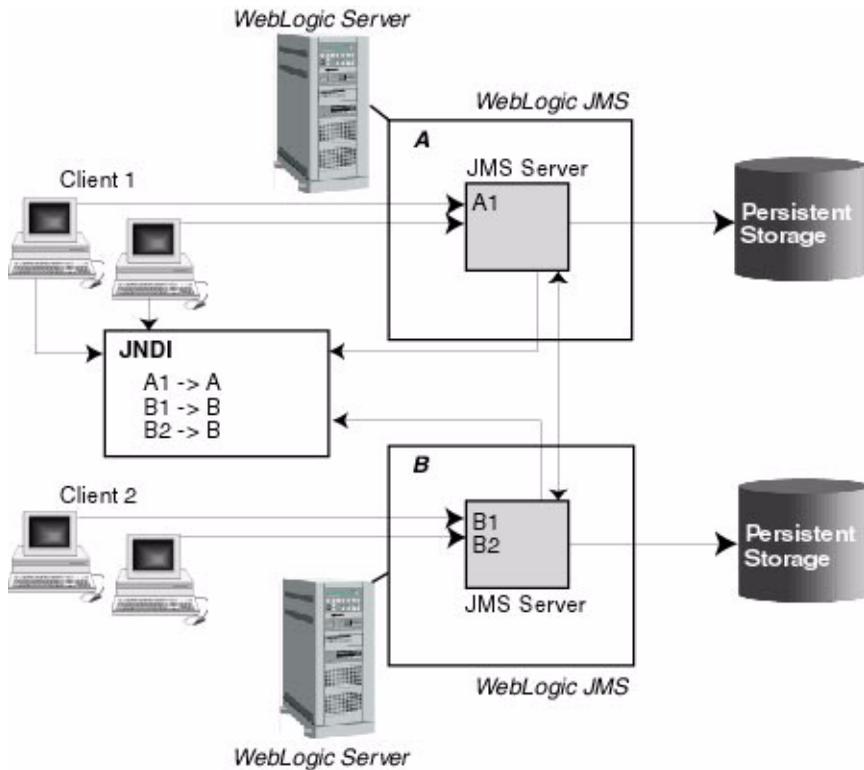
WebLogic JMS provides a full implementation of the JMS API. Specifically, WebLogic JMS:

- Provides a single, unified messaging API.
- Strictly adheres to the [JMS Specification – version 1.0.2b](#).
- Supports clustering.
- Supports messaging for applications that span different operating systems and machine architectures.
- Can be configured by setting attributes from the WebLogic Administration Console and/or using the JMS API to override values.
- Allows interoperability between JMS applications and other resource managers (primarily databases) using the Java Transaction API (JTA) transactions, including support for distributed transactions and the two-phase commit protocol. JMS applications can also participate in transactions with other Java APIs that use JTA, including non-WebLogic XA compliant message brokers.
- Supports messages containing Extensible Markup Language (XML).
- Supports multicasting allowing the delivery of messages to a select group of hosts using an IP multicast address.
- Can use either a database or a file for persistent message storage.
- Can be used with other WebLogic Server APIs and facilities, such as Enterprise Java Beans (EJB), JDBC connection pools, servlets, and RMI.

# WebLogic JMS Architecture

The following figure illustrates the WebLogic JMS architecture.

Figure 1-2 WebLogic JMS Architecture



## Major Components

The major components of the WebLogic JMS Server architecture, as illustrated in the figure “WebLogic JMS Architecture” on page 1-4, include:

- WebLogic JMS servers implementing the messaging facility
- Client applications
- JNDI (Java Naming and Directory Interface), which provides a server *lookup* facility

- Persistent storage (file or database) for storing persistent message data

## Clustering Features

The WebLogic JMS architecture implements *clustering* of multiple JMS servers by supporting cluster-wide, transparent access to destinations from any server in the cluster. Although WebLogic Server supports distributing JMS destinations and connection factories throughout a cluster, JMS topics and queues are still managed by individual WebLogic Server instances in the cluster.

For more information about configuring clustering for WebLogic JMS, see [“Configuring WebLogic JMS Clustering” on page 3-3](#). For detailed information about WebLogic Server clustering, see [Using WebLogic Server Clusters](#).

The advantages of clustering include the following:

- *Load balancing of destinations across multiple servers in the cluster*
  - An administrator can establish load balancing of destinations across multiple servers in the cluster by configuring multiple JMS servers and using *targets* to assign them to the defined WebLogic Servers. Each JMS server is deployed on exactly one WebLogic Server and handles requests for a set of destinations.

**Note:** Load balancing is not dynamic. During the configuration phase, the system administrator defines load balancing by specifying targets for JMS servers.

- An administrator can also configure multiple physical destinations as members of a single distributed destination set within a cluster. Producers and consumers are able to send and receive to the distributed destination. In the event of a single server failure within the cluster, WebLogic JMS then distributes the load across all available physical destination members within the distributed destination set.

For more information on distributed destinations, see [“Configuring Distributed Destinations”](#) in the *Administration Guide*.

- *Cluster-wide, transparent access to destinations from any server in the cluster*

A system administrator can establish cluster-wide, transparent access to destinations from any server in the cluster by either enabling the default connection factory for each server instance in the cluster, or by configuring one

or more connection factories and targeting them to one or more server instances in the cluster. This way, each connection factory can be deployed on multiple WebLogic Servers.

The application uses the Java Naming and Directory Interface (JNDI) to look up a connection factory and create a connection to establish communication with a JMS server. Each JMS server handles requests for a set of destinations. Requests for destinations not handled by a JMS server are forwarded to the appropriate server.

Connection factories are described in more detail in [“WebLogic JMS Fundamentals” on page 2-1](#).

## ■ *Scalability*

Scalability is provided by:

- Load balancing of destinations across multiple servers in the cluster, as described previously.
- Distribution of application load across multiple JMS servers via connection factories, thus reducing the load on any single JMS server and enabling session concentration by routing connections to specific servers.
- Optional multicast support, reducing the number of messages required to be delivered by a JMS server. The JMS server forwards only a single copy of a message to each host group associated with a multicast IP address, regardless of the number of applications that have subscribed.

## ■ *Migratability*

As an “exactly-once” service, WebLogic JMS takes advantage of the migration framework implemented in WebLogic Server for clustered environments. This allows WebLogic JMS to properly respond to migration requests and bring a JMS server online and offline in an orderly fashion. This includes both scheduled migrations as well as migrations in response to a WebLogic Server failure. For more information, see [“Configuring JMS Migratable Targets” on page 3-6](#).

**Note:** Automatic failover is not supported by WebLogic JMS for this release. For information about performing a manual failover, refer to [“Recovering from a WebLogic Server Failure” on page 3-11](#).

# WebLogic JMS Extensions

In addition to the API specified by Sun Microsystems' [JMS Specification](#), WebLogic JMS provides a public API, `weblogic.jms.extensions`, which includes classes and methods for the extensions described in the following table.

Extension	For more information. . .
Create XML messages	Refer to <a href="#">“Step 6a: Create the Message Object (Message Producers)”</a> on page 4-13
Define a session exception listener	Refer to <a href="#">“Defining a Session Exception Listener”</a> on page 4-48
Set or display the maximum number of pre-fetched asynchronous messages allowed on the session	Refer to <a href="#">“Dynamically Configuring Multicasting Configuration Attributes”</a> on page 4-85
Set or display the multicast session overrun policy that is applied when the message maximum is reached	Refer to <a href="#">“Dynamically Configuring Multicasting Configuration Attributes”</a> on page 4-85
Dynamically create permanent queues or topics	Refer to <a href="#">“Using the JMSHelper Class Methods”</a> on page 4-50
Convert between WebLogic JMS 7.0 and pre-6.0 <code>JMSMessageID</code> formats	Refer to <a href="#">“Setting Message Header Fields”</a> on page 4-59
Set a redelivery delay for messages	Refer to <a href="#">“Setting a Redelivery Delay for Messages”</a> on page 4-35
Set a message delivery time for producers	Refer to <a href="#">“Setting a Delivery Time on Producers”</a> on page 4-39
Set a delivery time for messages	Refer to <a href="#">“Setting a Delivery Time on Messages”</a> on page 4-39
Set a scheduled delivery time for messages	Refer to <a href="#">“Setting a Scheduled Time-to-Deliver Override”</a> on page 4-41

This API also supports `NO_ACKNOWLEDGE` and `MULTICAST_NO_ACKNOWLEDGE` acknowledge modes, and extended exceptions, including throwing an exception:

- To the session exception listener (if set), when one of its consumers has been closed by the server as a result of a server failure, or administrative intervention.

- From a multicast session when the number of messages received by the session but not yet delivered to the messages listener, exceeds the maximum number of messages allowed for that session.
- From a multicast consumer when it detects a sequence gap (message received out of sequence) in the data stream.

# JMS Enhancements in WebLogic Server 7.0

The following JMS enhancements are new to this release of WebLogic Server.

## High Availability Enhancement

WebLogic JMS takes advantage of the migration framework implemented in the WebLogic Server core for clustered environments. This allows WebLogic JMS to properly respond to migration requests and bring a JMS server online and offline in an orderly fashion. This includes both scheduled migrations as well as migrations in response to a WebLogic Server failure.

For more information, see “[Managing JMS](#)” in *Programming WebLogic JMS*.

## Distributed Destinations within a WebLogic Cluster

The highly available implementation of WebLogic JMS offers a level of service continuity in the event of a single server failure by enabling you to configure multiple physical destinations as members of a single distributed destination set. Specifically, an administrator can configure multiple instances of a given destination within a cluster. If one instance within the cluster fails, then other instances of the same destination will be able to provide service to JMS producers and consumers.

For more information, see “[Developing a WebLogic JMS Application](#)” in *Programming WebLogic JMS* and “[Managing JMS](#)” in the *Administration Guide*.

## Flow Control

Using the Flow Control feature, you can enable a JMS server or destination to slow down message producers when it determines that it is becoming overloaded. Specifically, when a JMS server/destination exceeds its specified bytes or messages thresholds, it instructs producers to limit their message flow.

For more information, see [“Managing JMS”](#) in the *Administration Guide*.

## WebLogic Messaging Bridge

A messaging bridge (also known as a *JMS bridge*) is responsible for transferring messages between two JMS providers. The WebLogic Server Messaging Bridge feature allows you to configure a store and forward mechanism between any two messaging providers—including separate implementations of WebLogic JMS.

For more information, see [“Using a WebLogic Messaging Bridge”](#) in the *Administration Guide*.

## Message Paging

The Message Paging feature can free up virtual memory during peak message load periods by swapping out messages from virtual memory to persistent storage when message loads reach a specified threshold. From a performance perspective, this feature can greatly benefit WebLogic Server implementations with the large message spaces that are required by today's enterprise applications.

For more information, see [“Managing JMS”](#) in the *Administration Guide*.



# 2 WebLogic JMS Fundamentals

The following sections describe WebLogic JMS components and features:

- [Messaging Models](#)
- [WebLogic JMS Classes](#)
- [ConnectionFactory](#)
- [Connection](#)
- [Session](#)
- [Destination](#)
- [Distributed Destination](#)
- [MessageProducer and MessageConsumer](#)
- [Message](#)
- [ServerSessionPoolFactory](#)
- [ServerSessionPool](#)
- [ServerSession](#)
- [ConnectionConsumer](#)

**Note:** For more information on the JMS classes described in this section, access the JMS Specification and Javadoc supplied on the Sun Microsystems' Java web site at the following location:

<http://java.sun.com/products/jms/docs.html>

# Messaging Models

JMS supports two messaging models: point-to-point (PTP) and publish/subscribe (Pub/sub). The messaging models are very similar, except for the following differences:

- PTP messaging model enables the delivery of a message to exactly one recipient.
- Pub/sub messaging model enables the delivery of a message to multiple recipients.

Each model is implemented with classes that extend common base classes. For example, the PTP class `javax.jms.Queue` and the Pub/sub class `javax.jms.Topic` both extend the class `javax.jms.Destination`.

Each message model is described in detail in the following sections.

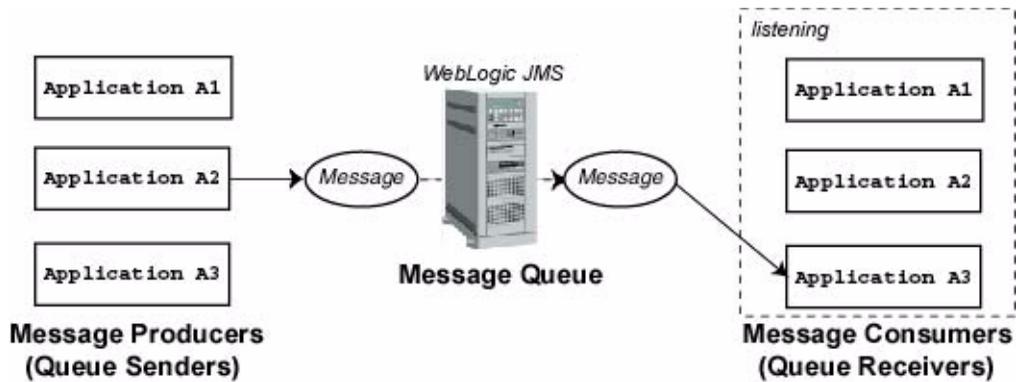
**Note:** The terms *producer* and *consumer* are used as generic descriptions of applications that send and receive messages, respectively, in either messaging model. For each specific messaging model, however, unique terms specific to that model are used when referring to producers and consumers.

## Point-to-Point Messaging

The point-to-point (PTP) messaging model enables one application to send a message to another. PTP messaging applications send and receive messages using named queues. A *queue sender* (producer) sends a message to a specific queue. A *queue receiver* (consumer) receives messages from a specific queue.

The following figure illustrates PTP messaging.

Figure 2-1 Point-to-Point (PTP) Messaging



Multiple queue senders and queue receivers can be associated with a single queue, but an individual message can be delivered to only *one* queue receiver.

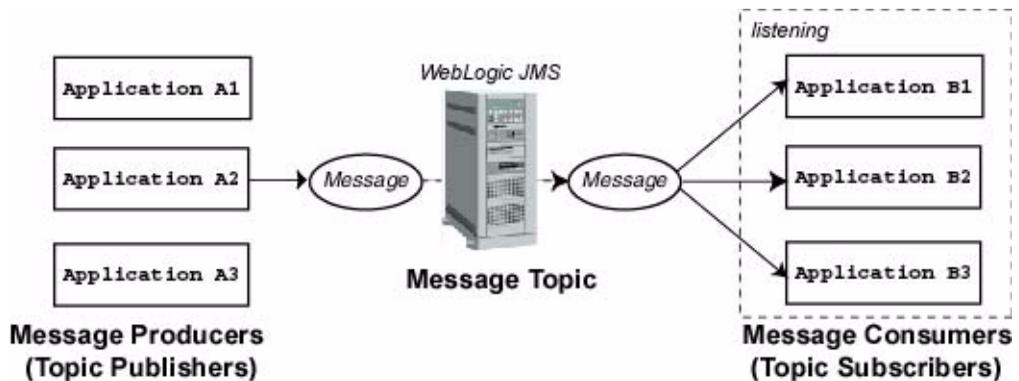
If multiple queue receivers are listening for messages on a queue, WebLogic JMS determines which one will receive the next message on a first come, first serve basis. If no queue receivers are listening on the queue, messages remain in the queue until a queue receiver attaches to the queue.

## Publish/Subscribe Messaging

The publish/subscribe (Pub/sub) messaging model enables an application to send a message to multiple applications. Pub/sub messaging applications send and receive messages by subscribing to a *topic*. A *topic publisher* (producer) sends messages to a specific topic. A *topic subscriber* (consumer) retrieves messages from a specific topic.

The following figure illustrates Pub/sub messaging.

Figure 2-2 Publish/Subscribe (Pub/Sub) Messaging



Unlike with the PTP messaging model, the Pub/sub messaging model allows multiple topic subscribers to receive the same message. JMS retains the message until all topic subscribers have received it.

The Pub/sub messaging model supports durable subscribers, allowing you to assign a name to a topic subscriber and associate it with a user or application. For more information about durable subscribers, see [“Setting Up Durable Subscriptions”](#) on page 4-53.

## Message Persistence

As per the “Message Delivery Mode” section of the [JMS Specification](#), messages can be specified as persistent or non-persistent:

- A persistent message is guaranteed to be delivered *once-and-only-once*. This means a message cannot be lost and cannot be delivered twice. It is not considered sent until it has been safely written to a file or database. WebLogic JMS writes persistent messages to a persistent backing store (disk-base file or JDBC-accessible database) assigned to each JMS server during configuration.
- Non-persistent messages are not stored. They are guaranteed to be delivered *at-most-once*, unless there is a system failure, in which case messages may be lost. If a connection is closed or recovered, all non-persistent messages that have

not yet been acknowledged will be redelivered. Once a non-persistent message is acknowledged, it will not be redelivered.

## WebLogic JMS Classes

To create a JMS applications, use the `javax.jms` API. The API allows you to create the class objects necessary to connect to the JMS, and send and receive messages. JMS class interfaces are created as subclasses to provide queue- and topic-specific versions of the common parent classes.

The following table lists the JMS classes described in more detail in subsequent sections. For a complete description of all JMS classes, see the `javax.jms`, `weblogic.jms.ServerSessionPoolFactory`, or `weblogic.jms.extensions` Javadoc.

**Table 2-1 WebLogic JMS Classes**

JMS Class	Description
<code>ConnectionFactory</code>	Encapsulates connection configuration information. A connection factory is used to create connections. You look up a connection factory using JNDI.
<code>Connection</code>	Represents an open communication channel to the messaging system. A connection is used to create sessions.
<code>Session</code>	Defines a serial order for the messages produced and consumed.
<code>Destination</code>	Identifies a queue or topic, encapsulating the address of a specific provider. Queue and topic destinations manage the messages delivered from the PTP and Pub/sub messaging models, respectively.
<code>MessageProducer</code> and <code>MessageConsumer</code>	Provides the interface for sending and receiving messages. Message producers send messages to a queue or topic. Message consumers receive messages from a queue or topic.
<code>Message</code>	Encapsulates information to be sent or received.

**Table 2-1 WebLogic JMS Classes**

JMS Class	Description
<code>ServerSessionPoolFactory</code> <sup>1</sup>	Encapsulates configuration information for a server-managed pool of message consumers. The server session pool factory is used to create server session pools.
<code>ServerSessionPool</code> <sup>1</sup>	Provides a pool of server sessions that can be used to process messages concurrently for connection consumers.
<code>ServerSession</code> <sup>1</sup>	Associates a thread with a JMS session.
<code>ConnectionConsumer</code> <sup>1</sup>	Specifies a consumer that retrieves server sessions to process messages concurrently.

<sup>1</sup> Supports an optional JMS interface for processing multiple messages concurrently.

For information about configuring JMS objects, see [“Managing WebLogic JMS” on page 3-1](#). The procedure for setting up a JMS application is presented in [“Setting Up a JMS Application” on page 4-4](#).

# ConnectionFactory

A `ConnectionFactory` object encapsulates connection configuration information, and enables JMS applications to create a `Connection`. A connection factory supports concurrent use, enabling multiple threads to access the object simultaneously. You can use the preconfigured default connection factory provided by WebLogic JMS, or a system administrator can configure one or more connection factories to create connections with predefined attributes that suit your application.

The default connection factory can be looked up using the JNDI name, `weblogic.jms.ConnectionFactory`. You only need to define a connection factory if the default provided by WebLogic JMS is not suitable for your application. All default factory configuration attributes are set to the same default values as a user-defined connection factory. For more information about the default values for the connection factory attributes, see [“JMS Connection Factory”](#) in the *Administration Console Online Help*.

Another distinction when using the default connection factory is that you have no control over targeting the WebLogic Server instances where the connection factory may be deployed. However, you can enable and/or disable the default connection factory on a per-server basis. For more information on enabling or disabling the default connection factory, see “[Server --> Services --> JMS](#)” in the *Administration Console Online Help*.

If the default connection factory is not suitable for your application, a system administrator can define and configure one or more connection factories to create connections with predefined attributes — as long as each connection factory is uniquely named — and WebLogic Server adds them to the JNDI space during startup. The application then retrieves a connection factory using WebLogic JNDI. For information on configuring and deploying user-defined connection factories, see “[Managing JMS](#)” in the *Administration Guide*.

A system administrator can also establish cluster-wide, transparent access to JMS destinations from any server in the cluster, either by enabling the default connection factory for each server instance in the cluster, or by configuring one or more connection factories and targeting them to one or more server instances in the cluster. This way, each connection factory can be deployed on multiple WebLogic Servers. For more information on JMS clustering, refer to “[Configuring WebLogic JMS Clustering](#)” on page 3-3.

**Notes:** For backwards compatibility, WebLogic JMS still supports two deprecated default connection factories. The JNDI names for these factories are:

```
javax.jms.QueueConnectionFactory and  
javax.jms.TopicConnectionFactory.
```

For information on migrating to a new default or user-defined connection factory from a deprecated connection factory, refer to “[Porting WebLogic JMS Applications](#)” on page 6-1.

The `ConnectionFactory` class does not define methods; however, its subclasses define methods for the respective messaging models. A connection factory supports concurrent use, enabling multiple threads to access the object simultaneously.

The following table describes the `ConnectionFactory` subclasses.

**Table 2-2 ConnectionFactory Subclasses**

<i>Subclass. . .</i>	<i>In Messaging Model. . .</i>	<i>Is Used to Create. . .</i>
<code>QueueConnectionFactory</code>	PTP	<code>QueueConnection</code> to a JMS PTP provider.
<code>TopicConnectionFactory</code>	Pub/sub	<code>TopicConnection</code> to a JMS Pub/sub provider.

To learn how to use the `ConnectionFactory` class within an application, see [“Developing a WebLogic JMS Application” on page 4-1](#), or the [`javax.jms.ConnectionFactory` Javadoc](#).

## Connection

A `Connection` object represents an open communication channel between an application and the messaging system, and is used to create a `Session` for producing and consuming messages. A connection creates server-side and client-side objects that manage the messaging activity between an application and JMS. A connection may also provide user authentication.

A `Connection` is created by a `ConnectionFactory`, obtained through a JNDI lookup.

Due to the resource overhead associated with authenticating users and setting up communications, most applications establish a single connection for all messaging. In the WebLogic Server, JMS traffic is multiplexed with other WebLogic services on the client connection to the server. No additional TCP/IP connections are created for JMS. Servlets and other server-side objects may also obtain JMS Connections.

By default, a connection is created in stopped mode. For information about how and when to start a stopped connection, see [“Starting, Stopping, and Closing a Connection” on page 4-46](#).

Connections support concurrent use, enabling multiple threads to access the object simultaneously.

The following table describes the `Connection` subclasses.

**Table 2-3 Connection Subclasses**

<i>Subclass. . .</i>	<i>In Messaging Model. . .</i>	<i>Is Used to Create. . .</i>
<code>QueueConnection</code>	PTP	<code>QueueSessions</code> , and consists of a connection to a JMS PTP provider created by <code>QueueConnectionFactory</code> .
<code>TopicConnection</code>	Pub/sub	<code>TopicSessions</code> , and consists of a connection to a JMS Pub/sub provider created by <code>TopicConnectionFactory</code> .

To learn how to use the `Connection` class within an application, see “[Developing a WebLogic JMS Application](#)” on page 4-1, or the `javax.jms.Connection` Javadoc.

## Session

A `Session` object defines a serial order for the messages produced and consumed, and can create multiple message producers and message consumers. The same thread can be used for producing and consuming messages. If an application wants to have a separate thread for producing and consuming messages, the application should create a separate session for each function.

A `Session` is created by the [Connection](#).

**Note:** A session and its message producers and consumers can only be accessed by one thread at a time. Their behavior is undefined if multiple threads access them simultaneously.

The following table describes the Session subclasses.

**Table 2-4 Session Subclasses**

<i>Subclass. . .</i>	<i>In Messaging Model. . .</i>	<i>Provides a Context for. . .</i>
QueueSession	PTP	Producing and consuming messages for a JMS PTP provider. Created by QueueConnection.
TopicSession	Pub/sub	Producing and consuming messages for a JMS Pub/sub provider. Created by TopicConnection.

To learn how to use the Session class within an application, see [“Developing a WebLogic JMS Application” on page 4-1](#), or the `javax.jms.Session` and `weblogic.jms.extensions.WLSession` javadocs.

## Non-transacted Session

In a non-transacted session, the application creating the session selects one of the five acknowledge modes defined in the following table.

**Table 2-5 Acknowledge Modes Used for Non-Transacted Sessions**

<b>Acknowledge Mode</b>	<b>Description</b>
AUTO_ACKNOWLEDGE	The Session object acknowledges receipt of a message once the receiving application method has returned from processing it.

**Table 2-5 Acknowledge Modes Used for Non-Transacted Sessions (Continued)**

Acknowledge Mode	Description
CLIENT_ACKNOWLEDGE	<p>The <code>Session</code> object relies on the application to call an <code>acknowledge</code> method on a received message. Once the method is called, the session acknowledges all messages received since the last <code>acknowledge</code>.</p> <p>This mode allows an application to receive, process, and acknowledge a batch of messages with one call.</p> <p><b>Note:</b> In the Administration Console, if the Acknowledge Policy attribute on the connection factory is set to <code>Previous</code>, but you want to acknowledge <i>all</i> received messages for a given session, then use the last message to invoke the <code>acknowledge</code> method. For more information on the Acknowledge Policy attribute, see “<a href="#">JMS Connection Factories</a>” in the <i>Administration Console Online Help</i>.</p>
DUPS_OK_ACKNOWLEDGE	<p>The <code>Session</code> object acknowledges receipt of a message once the receiving application method has returned from processing it; duplicate acknowledges are permitted.</p> <p>This mode is most efficient in terms of resource usage.</p> <p><b>Note:</b> You should avoid using this mode if your application cannot handle duplicate messages. Duplicate messages may be sent if an initial attempt to deliver a message fails.</p>
NO_ACKNOWLEDGE	<p>No acknowledge is required. Messages sent to a <code>NO_ACKNOWLEDGE</code> session are immediately deleted from the server. Messages received in this mode are not recovered, and as a result messages may be lost and/or duplicate message may be delivered if an initial attempt to deliver a message fails.</p> <p>This mode is supported for applications that do not require the quality of service provided by session acknowledge, and that do not want to incur the associated overhead.</p> <p><b>Note:</b> You should avoid using this mode if your application cannot handle lost or duplicate messages. Duplicate messages may be sent if an initial attempt to deliver a message fails.</p>

**Table 2-5 Acknowledge Modes Used for Non-Transacted Sessions (Continued)**

<b>Acknowledge Mode</b>	<b>Description</b>
MULTICAST_NO_ACKNOWLEDGE	<p>Multicast mode with no acknowledge required.</p> <p>Messages sent to a MULTICAST_NO_ACKNOWLEDGE session share the same characteristics as NO_ACKNOWLEDGE mode, described previously.</p> <p>This mode is supported for applications that want to support multicasting, and that do not require the quality of service provided by session acknowledge. For more information on multicasting, see <a href="#">“Using Multicasting” on page 4-81</a>.</p> <p><b>Note:</b> You should avoid using this mode if your application cannot handle lost or duplicate messages. Duplicate messages may be sent if an initial attempt to deliver a message fails.</p>

---

## Transacted Session

In a transacted session, only one transaction is active at any given time. Any messages sent or received during a transaction are treated as an atomic unit.

When you create a transacted session, the acknowledge mode is ignored. When an application commits a transaction, all the messages that the application received during the transaction are acknowledged by the messaging system and messages it sent are accepted for delivery. If an application rolls back a transaction, the messages that the application received during the transaction are not acknowledged and messages it sent are discarded.

JMS can participate in distributed transactions with other Java services, such as EJB, that use the Java Transaction API (JTA). Transacted sessions do not support this capability as the transaction is restricted to accessing the messages associated with that session. For more information about using JMS with JTA, see [“Using JTA User Transactions” on page 5-5](#).

# Destination

A `Destination` object can be either a queue or topic, encapsulating the address syntax for a specific provider. The JMS specification does not define a standard address syntax due to the variations in syntax between providers.

Similar to a connection factory, an administrator defines and configures the destination and the WebLogic Server adds it to the JNDI space during startup. Applications can also create temporary destinations that exist only for the duration of the JMS connection in which they are created.

**Note:** Administrators can also configure multiple physical destinations as members of a single distributed destination set within a server cluster. For more information, see [“Distributed Destination”](#) on page 2-14.

On the client side, `Queue` and `Topic` objects are handles to the object on the server. Their methods only return their names. To access them for messaging, you create message producers and consumers that attach to them.

A destination supports concurrent use, enabling multiple threads to access the object simultaneously. JMS `Queues` and `Topics` extend `javax.jms.Destination`. The following table describes the `Destination` subclasses.

**Table 2-6 Destination Subclasses**

<i>Subclass . .</i>	<i>In Messaging Model . .</i>	<i>Manages Messages for . .</i>
<code>Queue</code>	PTP	JMS PTP provider.
<code>TemporaryQueue</code>	PTP	JMS PTP provider, and exists for the duration of the JMS connection in which the messages are created. A temporary queue can be consumed only by the queue connection that created it.
<code>Topic</code>	Pub/sub	JMS Pub/sub provider.
<code>TemporaryTopic</code>	Pub/sub	JMS Pub/sub provider, and exists for the duration of the JMS connection in which the messages are created. A temporary topic can be consumed only by the topic connection that created it.

**Note:** An application has the option of browsing queues by creating a `QueueBrowser` object in its queue session. This object produces a *snapshot* of the messages in the queue at the time the queue browser is created. The application can view the messages in the queue, but the messages are not considered *read* and are not removed from the queue. For more information about browsing queues, see [“Browsing Header and Property Fields” on page 4-65](#).

To learn how to use the `Destination` class within an application, see [“Developing a WebLogic JMS Application” on page 4-1](#), or the `javax.jms.Destination` Javadoc.

# Distributed Destination

Administrators can configure multiple physical destinations as members of a single distributed destination set within a WebLogic Server cluster. Once properly configured, your producers and consumers are able to send and receive to the distributed destination. WebLogic JMS then distributes the messaging load across all available destination members within the distributed destination.

- For more information on using distributed destinations with your applications, see [“Using Distributed Destinations” on page 4-88](#).
- For instructions on configuring distributed destinations using the Administration Console, see [“Configuring Distributed Destinations”](#) in the *Administration Guide*.

# MessageProducer and MessageConsumer

A `MessageProducer` object sends messages to a queue or topic. A `MessageConsumer` object receives messages from a queue or topic. Message producers and consumers operate independently of one another. Message producers generate and send messages regardless of whether a message consumer has been created and is waiting for a message, and vice versa.

A `Session` creates the `MessageProducers` and `MessageConsumers` that are attached to queues and topics.

The message sender and receiver objects are created as subclasses of the `MessageProducer` and `MessageConsumer` classes. The following table describes the `MessageProducer` and `MessageConsumer` subclasses.

**Table 2-7 MessageProducer and MessageConsumer Subclasses**

<i>Subclass. . .</i>	<i>In Messaging Model. . .</i>	<i>Performs the Following Function. . .</i>
<code>QueueSender</code>	PTP	Sends messages for a JMS PTP provider.
<code>QueueReceiver</code>	PTP	Receives messages for a JMS PTP provider, and exists until the JMS connection in which the messages are created is closed.
<code>TopicPublisher</code>	Pub/sub	Sends messages for a JMS Pub/sub provider.
<code>TopicSubscriber</code>	Pub/sub	Receives messages for a JMS Pub/sub provider, and exists for the duration of the JMS connection in which the messages are created. Message destinations must be bound explicitly using the appropriate JNDI interface.

The PTP model, as shown in the figure “[Point-to-Point \(PTP\) Messaging](#)” on page 2-3, allows multiple sessions to receive messages from the same queue. However, a message can only be delivered to one queue receiver. When there are multiple queue receivers, WebLogic JMS defines the next queue receiver that will receive a message on a first-come, first-serve basis.

The Pub/sub model, as shown in the figure “[Publish/Subscribe \(Pub/Sub\) Messaging](#)” on page 2-4, allows messages to be delivered to multiple topic subscribers. Topic subscribers can be durable or non-durable, as described in “[Setting Up Durable Subscriptions](#)” on page 4-53.

An application can use the same JMS connection to both publish and subscribe to a single topic. Because topic messages are delivered to all subscribers, an application can receive messages it has published itself. To prevent clients from receiving messages that they publish, a JMS application can set a `noLocal` attribute on the topic subscriber, as described in “[Step 5: Create Message Producers and Message Consumers Using the Session and Destinations](#)” on page 4-11.

To learn how to use the `MessageProducer` and `MessageConsumer` classes within an application, see [“Setting Up a JMS Application” on page 4-4](#), or the [`javax.jms.MessageProducer` and `javax.jms.MessageConsumer` javadocs](#).

# Message

A `Message` object encapsulates the information exchanged by applications. This information includes three components: a set of standard header fields, a set of application-specific properties, and a message body. The following sections describe these components.

## Message Header Fields

Every JMS message contains a standard set of header fields that is included by default and available to message consumers. Some fields can be set by the message producers.

For information about setting message header fields, see [“Setting and Browsing Message Header and Property Fields” on page 4-58](#), or to the [`javax.jms.Message` Javadoc](#).

The following table describes the fields in the message headers and shows how values are defined for each field.

**Table 2-8 Message Header Fields**

Field	Description	Defined by
JMSCorrelationID	<p>Specifies one of the following: a WebLogic JMSMessageID (described later in this table), an application-specific string, or a byte[] array. The JMSCorrelationID is used to correlate messages.</p> <p>There are two common applications for this field.</p> <p>The first application is to link messages by setting up a request/response scheme, as follows:</p> <ol style="list-style-type: none"> <li>1. When an application sends a message, it stores the JMSMessageID value assigned to it.</li> <li>2. When an application receives the message, it copies the JMSMessageID into the JMSCorrelationID field of a response message that it sends back to the sending application.</li> </ol> <p>The second application is to use the JMSCorrelationID field to carry any String you choose, enabling a series of messages to be linked with some application-determined value.</p> <p>All JMSMessageIDs start with an ID: prefix. If you use the JMSCorrelationID for some other application-specific string, it <i>must not</i> begin with the ID: prefix.</p>	Application

**Table 2-8 Message Header Fields (Continued)**

<b>Field</b>	<b>Description</b>	<b>Defined by</b>
<code>JMSDeliveryMode</code>	<p>Specifies <code>PERSISTENT</code> or <code>NON_PERSISTENT</code> messaging.</p> <p>When a persistent message is sent, WebLogic JMS stores it in the JMS file or JDBC database. The <code>send()</code> operation is not considered successful until delivery of the message can be guaranteed. A persistent message is guaranteed to be delivered at least once.</p> <p>WebLogic JMS does not store non-persistent messages in the JMS database. This mode of operation provides the lowest overhead. They are guaranteed to be delivered at least once unless there is a system failure, in which case messages may be lost. If a connection is closed or recovered, all non-persistent messages that have not yet been acknowledged will be redelivered. Once a non-persistent message is acknowledged, it will not be redelivered.</p> <p>When a message is sent, this value is ignored. When the message is received, it contains the delivery mode specified by the sending method.</p>	<code>send()</code> method
<code>JMSDeliveryTime</code>	<p>Defines the earliest absolute time at which a message can be delivered to a consumer. This field can be used to sort messages in a destination and to select messages. For purposes of data type conversion, the <code>JMSDeliveryTime</code> is a long integer.</p>	<code>send()</code> method
<code>JMSDestination</code>	<p>Specifies the destination (queue or topic) to which the message is to be delivered. The application's message producer sets the value of this field when the message is sent.</p> <p>When a message is sent, this value is ignored. When a message is received, its destination value must be equivalent to the value assigned when it was sent.</p>	<code>send()</code> method
<code>JMSExpiration</code>	<p>Specifies the expiration, or time-to-live value, for a message.</p> <p>WebLogic JMS calculates the <code>JMSExpiration</code> value as the sum of the application's time-to-live and the current GMT. If the application specifies time-to-live as 0, <code>JMSExpiration</code> is set to 0, which means the message never expires.</p> <p>WebLogic JMS removes expired messages from the system to prevent their delivery.</p>	<code>send()</code> method

**Table 2-8 Message Header Fields (Continued)**

Field	Description	Defined by
<code>JMSMessageID</code>	<p>Contains a string value that uniquely identifies each message sent by a JMS Provider.</p> <p>All <code>JMSMessageID</code>s start with an <code>ID:</code> prefix.</p> <p>When a message is sent, this value is ignored. When the message is received, it contains a provider-assigned value.</p>	<code>send()</code> method
<code>JMSPriority</code>	<p>Specifies the priority level. This field is set before a message is sent.</p> <p>JMS defines ten priority levels, 0 to 9, 0 being the lowest priority. Levels 0-4 indicate gradations of <i>normal</i> priority, and level 5-9 indicate gradations of <i>expedited</i> priority.</p> <p>When the message is received, it contains the value specified by the method sending the message.</p> <p>You can sort destinations by priority by configuring a destination key, as described in <a href="#">Managing JMS</a> in the <i>Administration Guide</i>.</p>	<code>send()</code> method
<code>JMSRedelivered</code>	<p>Specifies a flag set when a message is redelivered because no acknowledge was received. This flag is of interest to a receiving application only.</p> <p>If set, the flag indicates that JMS may have delivered the message previously because one of the following is true:</p> <ul style="list-style-type: none"> <li>■ The application has already received the message, but did not acknowledge it.</li> <li>■ The session's <code>recover()</code> method was called to restart the session beginning after the last acknowledged message. For more information about the <code>recover()</code> method, see <a href="#">“Recovering Received Messages” on page 4-33</a>.</li> </ul>	WebLogic JMS

**Table 2-8 Message Header Fields (Continued)**

Field	Description	Defined by
JMSReplyTo	<p>Specifies a queue or topic to which reply messages should be sent. This field is set by the sending application before the message is sent.</p> <p>This feature can be used with the <code>JMSCorrelationID</code> header field to coordinate request/response messages.</p> <p>Simply setting the <code>JMSReplyTo</code> field does not guarantee a response; it <i>enables</i> the receiving application to respond, if it so chooses.</p> <p>You may set the <code>JMSReplyTo</code> to null, which may have a semantic meaning to the receiving application, such as a notification event.</p>	Application
JMSTimestamp	<p>Contains the time at which the message was sent. WebLogic JMS writes the timestamp in the message when it accepts the message for delivery, <i>not</i> when the application sends the message.</p> <p>When the message is received, it contains the timestamp.</p> <p>The value stored in the field is a Java millis time value.</p>	WebLogic JMS
JMSType	<p>Specifies the message type identifier (String) set by the sending application.</p> <p>The JMS specification allows some flexibility with this field in order to accommodate diverse JMS providers. Some messaging systems allow application-specific message types to be used. For such systems, the <code>JMSType</code> field could be used to hold a message type ID that provides access to the stored type definitions.</p> <p>WebLogic JMS does not restrict the use of this field.</p>	Application

## Message Property Fields

The property fields of a message contain header fields added by the sending application. The properties are standard Java name/value pairs. Property names must conform to the message selector syntax specifications defined in the [javax.jms.Message](#) Javadoc. The following values are valid: boolean, byte, double, float, int, long, short, and String.

Although message property fields may be used for application-specific purposes, JMS provides them primarily for use in message selectors. For more information about message selectors, see [“Filtering Messages” on page 4-66](#).

For information about setting message property fields, see [“Setting and Browsing Message Header and Property Fields” on page 4-58](#), or to the `javax.jms.Message` Javadoc.

## Message Body

A message body contains the content being delivered from producer to consumer.

The following table describes the types of messages defined by JMS. All message types extend `javax.jms.Message`, which consists of message headers and properties, but no message body.

**Table 2-9 JMS Message Types**

Type	Description
<code>javax.jms.BytesMessage</code>	Stream of uninterpreted bytes, which must be understood by the sender and receiver. The access methods for this message type are stream-oriented readers and writers based on <code>java.io.DataInputStream</code> and <code>java.io.DataOutputStream</code> .
<code>javax.jms.MapMessage</code>	Set of name/value pairs in which the names are strings and the values are Java primitive types. Pairs can be read sequentially or randomly, by specifying a name.
<code>javax.jms.ObjectMessage</code>	Single serializable Java object.
<code>javax.jms.StreamMessage</code>	Similar to a <code>BytesMessage</code> , except that only Java primitive types are written to or read from the stream.
<code>javax.jms.TextMessage</code>	Single String. The <code>TextMessage</code> can also contain XML content.
<code>weblogic.jms.extensions.XMLMessage</code>	XML content. Use of the <code>XMLMessage</code> type facilitates message filtering, which is more complex when performed on XML content shipped in a <code>TextMessage</code> .

For more information, see the `javax.jms.Message` Javadoc. For more information about the access methods and, if applicable, the conversion charts associated with a particular message type, see the Javadoc for that message type.

# ServerSessionPoolFactory

A server session pool is a WebLogic-specific JMS feature that enables you to process messages concurrently. A server session pool factory is used to create a server-side `ServerSessionPool`.

WebLogic JMS defines one `ServerSessionPoolFactory` object, by default: `weblogic.jms.ServerSessionPoolFactory:<name>`, where `<name>` specifies the name of the JMS server to which the session pool is created. The WebLogic Server adds the default server session pool factory to the JNDI space during startup and the application subsequently retrieves the server session pool factory using WebLogic JNDI.

To learn how to use the server session pool factory within an application, see “[Defining Server Session Pools](#)” on page 4-71, or the `weblogic.jms.ServerSessionPoolFactory` Javadoc.

# ServerSessionPool

A `ServerSessionPool` application server object provides a pool of server sessions that connection consumers can retrieve in order to process messages concurrently.

A `ServerSessionPool` is created by the `ServerSessionPoolFactory` object obtained through a JNDI lookup.

To learn how to use the server session pool within an application, see “[Defining Server Session Pools](#)” on page 4-71, or the `javax.jms.ServerSessionPool` Javadoc.

# ServerSession

A `ServerSession` application server object enables you to associate a thread with a JMS session by providing a context for creating, sending, and receiving messages.

A `ServerSession` is created by a `ServerSessionPool` object.

To learn how to use the server session within an application, see “[Defining Server Session Pools](#)” on page 4-71, or the `javax.jms.ServerSession` Javadoc.

# ConnectionConsumer

A `ConnectionConsumer` object uses a server session to process received messages. If message traffic is heavy, the connection consumer can load each server session with multiple messages to minimize thread context switching.

A `ConnectionConsumer` is created by a `Connection` object.

To learn how to use the connection consumers within an application, see “[Defining Server Session Pools](#)” on page 4-71, or the `javax.jms.ConnectionConsumer` Javadoc.

**Note:** Connection consumer listeners run on the same JVM as the server.



# 3 Managing WebLogic JMS

The *Administration Console* provides the interface that you can use to enable, configure, and monitor the features of the WebLogic Server, including JMS. To invoke the *Administration Console*, refer to the procedures described in “[Starting and Using the Administration Console](#)” in the *Administration Guide*.

The following sections provide an overview of configuring and monitoring WebLogic JMS:

- [Configuring WebLogic JMS](#)
- [Configuring WebLogic JMS Clustering](#)
- [Configuring JMS Migratable Targets](#)
- [Tuning WebLogic JMS](#)
- [Monitoring WebLogic JMS](#)
- [Recovering from a WebLogic Server Failure](#)

## Configuring WebLogic JMS

Using the Administration Console, you define configuration attributes to:

- Enable JMS.
- Create JMS servers.

- Create and/or customize values for JMS servers, connection factories, destinations (physical queues and topics), distributed destinations (sets of physical queue and topic members within a cluster) destination templates, destination sort ordering (using destination keys), persistent stores, paging stores, session pools, and connection consumers.
- Set up custom JMS applications.
- Define thresholds and quotas.
- Enable any desired JMS features, such as:
  - server clustering
  - concurrent message processing
  - persistent messaging
  - message paging
  - flow control

WebLogic JMS provides default values for some configuration attributes; you must provide values for all others. If you specify an invalid value for any configuration attribute, or if you fail to specify a value for an attribute for which a default does not exist, the WebLogic Server will not boot JMS when you restart it. A sample JMS configuration is provided with the product.

When migrating from a previous release, the configuration information will be converted automatically, as described in [“Porting Existing Applications” on page 6-8](#).

**Note:** [Appendix A, “Configuration Checklists,”](#) provides checklists that enable you to view the attribute requirements and/or options for supporting various JMS features.

# Configuring WebLogic JMS Clustering

A WebLogic Server *cluster* is a group of servers that work together to provide a more scalable, more reliable application platform than a single server. A cluster appears to its clients as a single server but is in fact a group of servers acting as one. A cluster provides three key features above a single server:

- Scalability—servers can be added to the cluster dynamically to increase capacity.
- High Availability—redundancy of multiple servers insulates applications from failures. Redundancy of multiple destinations (queues and topics) as members of a single *distributed destination* set within a cluster ensures redistribution of the messaging load to other available members in the set when one member becomes unavailable.
- Migratability—respond to migration requests and bring a JMS server online and offline in an orderly fashion. This includes both scheduled migrations as well as migrations in response to a WebLogic Server failure.

**Note:** JMS clients depend on unique WebLogic Server names to successfully access a cluster—even when WebLogic Servers reside in different domains. Therefore, make sure that *all* WebLogic Servers that JMS clients contact have unique server names.

For more information about starting WebLogic clusters and its features and benefits, see “[Configuring WebLogic Servers and Clusters](#)” in *Using WebLogic Server Clusters*.

## How JMS Clustering Works

An administrator can establish cluster-wide, transparent access to JMS destinations from any server in the cluster, either by using the default connection factory for each server instance in the cluster, or by configuring one or more connection factories and targeting them to one or more server instances in the cluster. This way, each connection factory can be deployed on multiple WebLogic Servers. However, each user-defined connection factory must be uniquely named to be successfully deployed on multiple WebLogic Servers.

The administrator can configure multiple JMS servers on the various nodes in the cluster—as long as the JMS servers are uniquely named—and can then assign JMS destinations to the various JMS servers.

The application uses the Java Naming and Directory Interface (JNDI) to look up a connection factory and create a connection to establish communication with a JMS server. Each JMS server handles requests for a set of destinations. Requests for destinations not handled by a JMS server are forwarded to the appropriate WebLogic Server.

### **JMS Clustering Naming Requirements**

JMS clients depend on unique WebLogic Server names to successfully access a cluster—even when WebLogic Servers reside in different domains. Therefore, make sure that all WebLogic Servers that JMS clients contact have unique server names.

The following unique naming requirements apply when configuring WebLogic JMS to work in a clustered environment in either a single WebLogic domain or in a multi-domain environment.

- All JMS connection factories targeted to nodes in a domain must be uniquely named.
- All JMS servers targeted to nodes in all domains must be uniquely named.
- If persistent messaging is required, all JMS stores in all domains must be uniquely named.

For more information about JMS resource naming rules in single domain and multi-domain environments, see “[JMS Resource Naming Rules for Domain Interoperability](#)” in the *Administration Guide*.

### **JMS Distributed Destination within a Cluster**

An administrator can also configure multiple JMS destinations in cluster as part of a single distributed destination, which is a set of queues or topics that are called under a single JNDI name so This way they appear to be a single, logical destination to a JMS client, when the members of the set are actually distributed across multiple servers within a cluster, with each destination member belonging to a separate JMS server. This way, in the event of a single server failure within the cluster, the message load is

distributed across the other available members of the distributed destination. For more information on configuring a distributed destination, see “[Configuring Distributed Destinations](#)” in the *Administration Guide*.

### JMS as a Migratable Service within a Cluster

WebLogic JMS takes advantage of the migration framework implemented in the WebLogic Server core for clustered environments. This allows WebLogic JMS to properly respond to migration requests and bring a JMS server online and offline in an orderly fashion. This includes both scheduled migrations as well as migrations in response to a WebLogic Server failure. For more information, see “[Configuring JMS Migratable Targets](#)” on page 3-6.

## Configuration Guidelines for JMS Clustering

In order to use WebLogic JMS in a clustered environment, follow these guidelines:

1. Configure your clustered environment as described in “[Setting Up WebLogic Clusters](#)” in *Using WebLogic Server Clusters*.
2. Identify server targets for JMS connection factories using the *Administration Console*. For connection factories, you can identify either a single-server target or a cluster target, which are server instances that are associated with a connection factory to support clustering.

For more information about these connection factory configuration attributes, see “[Configuring Connection Factories](#)” in the *Administration Guide*.

3. Optionally, identify migratable server targets for JMS servers using the *Administration Console*. For JMS servers, you can identify either a single-server target or a migratable target, which is a set of server instances in a cluster that can host an “exactly-once” service like JMS in case of a single server failure.

For more information on migratable JMS server targets, see “[Configuring JMS Migratable Targets](#)” on page 3-6. For more information about JMS server configuration attributes, see “[Configuring JMS Servers](#)” in the *Administration Guide*.

**Note:** You cannot deploy the same destination on more than one JMS server. In addition, you cannot deploy a JMS server on more than one WebLogic Server.

4. Optionally, you can configure the physical JMS destinations in a cluster as part of a distributed destination set, as discussed in [“JMS Distributed Destination within a Cluster”](#) on page 3-4.

## What About Failover?

For WebLogic JMS implementations that are part of a WebLogic 7.0 clustered environment, JMS offers service continuity in the event of a single Weblogic Server failure by enabling you to configure multiple physical destinations (queues and topics) as part of a single distributed destination set. In addition, implementing the Migratable Service feature, will ensure that pinned “exactly-once” services, like JMS, do not introduce a single point of failure for dependent applications in the cluster,

However, automatic failover is not currently supported by WebLogic JMS. For information about performing a manual failover, refer to [“Recovering from a WebLogic Server Failure”](#) on page 3-11.

## Configuring JMS Migratable Targets

As an exactly-once service, WebLogic JMS is not active on all WebLogic Server instances in a cluster. It is instead “pinned” to a single server in the cluster to preserve data consistency. To ensure that pinned services do not introduce a single point of failure for dependent applications in the cluster, WebLogic Server can be configured to migrate exactly-once services to any server in the migratable target list.

WebLogic JMS takes advantage of the migration framework by allowing an administrator to specify a migratable target for a JMS server in the *Administration Console*. Once properly configured, a JMS server and all of its destinations can migrate to another WebLogic Server within a cluster.

This allows WebLogic JMS to properly respond to migration requests and bring a JMS server online and offline in an orderly fashion. This includes both scheduled migrations as well as migrations in response to a WebLogic Server failure with the cluster.

For more information about defining migratable targets, see “[Migration for Pinned Services](#)” in *Using WebLogic Server Clusters*.

## How JMS Migration Works

For implementations that are part of a WebLogic clustered environment, WebLogic JMS implements the server migration framework, which allows JMS servers to respond to activate and deactivate requests.

**Table 3-1 WebLogic JMS Migration Process**

Migration state...	What takes place...
Initialization	Initialization of a JMS server includes processing any configuration or deployment information and creating the appropriate objects. Destinations and other JMS resources are unavailable at this time. In addition, the persistent store is not opened, as this could compromise the integrity of the store. The JMS server makes itself available to handle changes in configuration that may occur between initialization and activation.
Activation	When a JMS server is activated, it opens the persistent store, performs any necessary recovery, reconciles the contents of the store with the current configuration, and makes the destinations available for access by applications. In addition, any configured server session pools begin processing after activation is complete.
Deactivation	When a JMS server is deactivated it stops all server session pool processing, marks all destinations as unavailable, flushes and closes its persistent stores, purges its destinations, and deletes all objects for the JMS server.

## Configuration Steps for JMS Server Migration

In order to make a JMS server a migratable service in a clustered environment, you must do the following:

1. Optionally, familiarize yourself with how server migration for pinned services works by reading “[Migration for Pinned Services](#)” in *Using WebLogic Server Clusters*.
2. For JMS implementations that use persistent messaging, make sure that the JMS store is configured such that all the candidate servers in a migratable target share access to the store. For more information about migrating JMS stores, see “[Persistent Store Migration](#)” on page 3-8.
3. Configure a migratable target server for the cluster that can potentially host a JMS server, as described in “[Configure Migratable Targets for Pinned Services](#)” in *Using WebLogic Server Clusters*.  
**Note:** You must set a unique Listen Address value for the migratable target server instance that will host a migrated the JMS server; otherwise, the migration will fail.
4. Identify a migratable target server instance on which to deploy a JMS server as described in “[Deploying JMS to a Migratable Target Server Instance](#)” in *Using WebLogic Server Clusters*.  
**Note:** When a migratable target server boots, the JMS server automatically boots as well on the user-preferred server in the cluster.
5. You can also manually migrate a JMS server and all of its destinations before performing server maintenance or to a healthy server if the host server fails. For more information, see “[Migrating a Pinned Service To a Target Service Instance](#)” in *Using WebLogic Server Clusters*.  
**Note:** A JMS server’s distributed destination members can migrate to another server instance within a cluster—even when the target server instance is already hosting a JMS server with its own distributed destination members. For more information about distributed destination failover, see “[Distributed Destination Failover](#)” on page 4-102.

### **Persistent Store Migration**

Weblogic JMS persistent stores cannot be migrated along with JMS servers; therefore, applications that need access to persistent stores from other physical machines after the migration of a JMS server must implement an alternative solution, as follows:

- Implement a hardware solution, such as a dual-ported SCSI disk or Storage Area Network (SAN) to make your JMS persistent store available from other machines.
- Use JDBC to access your JMS JDBC store, which can be on yet another server. Applications can then take advantage of any high-availability or failover solutions offered by your database vendor.

For more information about configuring a JMS JDBC store, see “[Configuring Stores](#)” in the *Administration Guide*.

## Migration Failover

For information about procedures for recovering from a WebLogic Server failure, see “[Managing JMS](#)” in the *Administration Guide*.

# Tuning WebLogic JMS

The following sections explain how to get the most out of your applications by implementing the administrative performance tuning features available with WebLogic JMS.

- Synchronous Write Policies for JMS File Stores—disabling synchronous writes improves file store performance, often quite dramatically, but at the expense of possibly losing sent messages or generating duplicate received messages in the event of an operating system crash or a hardware failure.

For more information, see “[Configuring a Synchronous Write Policy for JMS File Stores](#)” in the *Administration Guide*.

- Using Message Paging—you can free up valuable virtual memory during peak message load periods by swapping out messages from memory to persistent storage whenever your message loads reach a specified threshold. From a performance perspective, this feature can greatly benefit WebLogic Server implementations with the large message spaces that are required by today's enterprise applications.

For more information, see “[Using Message Paging](#)” in the *Administration Guide*.

- Establishing Message Flow Control—a JMS server or JMS destination (queue or topic) can be configured to instruct message producers to limit their message flow when it determines that it is becoming overloaded.

For more information, see “[Establishing Message Flow Control](#)” in the *Administration Guide*.

- Tuning Distributed Destinations—the following attributes on the JMS Connection Factory can be configured to tune your distributed destinations:
  - Load Balancing—defines whether WebLogic JMS will spread or balance the messaging load across distributed destinations.
  - Server Affinity—defines whether a WebLogic Server that is attempting to load balance consumers or producers across multiple physical destinations in a distributed destination set, will first attempt to load balance across those physical destinations being served by any JMS servers that are also running on the same WebLogic Server.

For more information, see “[Tuning Distributed Destinations](#)” in the *Administration Guide*.

# Monitoring WebLogic JMS

Statistics are provided for the following JMS objects: JMS servers, connections, sessions, destinations, durable subscribers, message producers, message consumers, and server session pools. You can monitor JMS statistics using the *Administration Console*.

JMS statistics continue to increment as long as the server is running. Statistics can only be reset when the server is rebooted.

For more information on configuring and monitoring WebLogic JMS, see “[Managing JMS](#)” in the *Administration Guide*.

Once WebLogic JMS has been configured, applications can begin sending and receiving messages through the JMS API, as described in “[Developing a WebLogic JMS Application](#)” on page 4-1.

# Recovering from a WebLogic Server Failure

The procedures for recovering from a WebLogic Server failure, and performing a manual failover, including programming considerations, are described in detail in “[Managing JMS](#)” in the *Administration Guide*.



# 4 Developing a WebLogic JMS Application

The following sections describe how to develop a WebLogic JMS application:

- [Application Development Flow](#)
- [Importing Required Packages](#)
- [Setting Up a JMS Application](#)
- [Sending Messages](#)
- [Receiving Messages](#)
- [Acknowledging Received Messages](#)
- [Releasing Object Resources](#)
- [Managing Rolled Back, Recovered, or Expired Messages](#)
- [Setting Message Delivery Times](#)
- [Managing Connections](#)
- [Managing Sessions](#)
- [Using Temporary Destinations](#)
- [Setting Up Durable Subscriptions](#)
- [Setting and Browsing Message Header and Property Fields](#)
- [Filtering Messages](#)
- [Defining Server Session Pools](#)

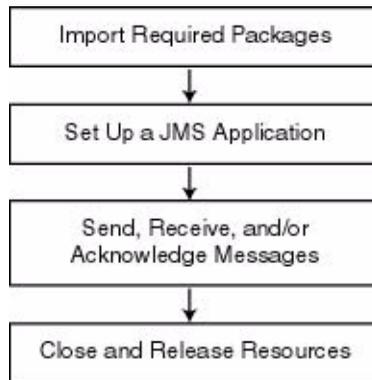
- [Using Multicasting](#)
- [Using Distributed Destinations](#)

**Note:** For more information about the JMS classes described in this section, access the JMS Javadoc supplied on the Sun Microsystems' Java web site at the following location: <http://java.sun.com/products/jms/docs.html>

# Application Development Flow

When developing a WebLogic JMS application, you must perform the steps identified in the following figure.

**Figure 4-1 WebLogic JMS Application Development Flow—Required Steps**



In addition to the application development steps defined in the previous figure, you can also optionally perform any of the following steps during your design development:

- Manage connection and session processing
- Create destinations dynamically
- Create durable subscriptions
- Manage message processing by setting and browsing message header and property fields, filtering messages, and/or processing messages concurrently

- Use multicasting
- Use JMS within transactions (described in [“Using Transactions with WebLogic JMS”](#) on page 5-1)

Except where noted, all application development steps are described in the following sections.

## Importing Required Packages

The following table lists the packages that are commonly used by WebLogic JMS applications.

**Table 4-1 WebLogic JMS Packages**

Package	Description
<a href="#">javax.jms</a>	Sun Microsystems’ JMS API. This package is always used by WebLogic JMS applications.
<a href="#">java.util</a>	Utility API, such as date and time facilities.
<a href="#">java.io</a>	System input and output API.
<a href="#">javax.naming</a> <a href="#">weblogic.jndi</a>	JNDI packages required for server and destination lookups.
<a href="#">javax.transaction.UserTransaction</a>	JTA API required for JTA user transaction support.
<a href="#">weblogic.jms.ServerSessionPoolFactory</a>	WebLogic JMS public API for use with server session pools, an optional application server facility described in the JMS specification.
<a href="#">weblogic.jms.extensions</a>	WebLogic-specific JMS public API that provides additional classes and methods, as described in <a href="#">“WebLogic JMS Extensions”</a> on page 1-7.

Include the following package `import` statements at the beginning of your program:

```
import javax.jms.*;
import java.util.*;
```

```
import java.io.*;
import javax.naming.*;
import javax.transaction.*;
```

If you implement a server session pool application, also include the following class on your import list:

```
import weblogic.jms.ServerSessionPoolFactory;
```

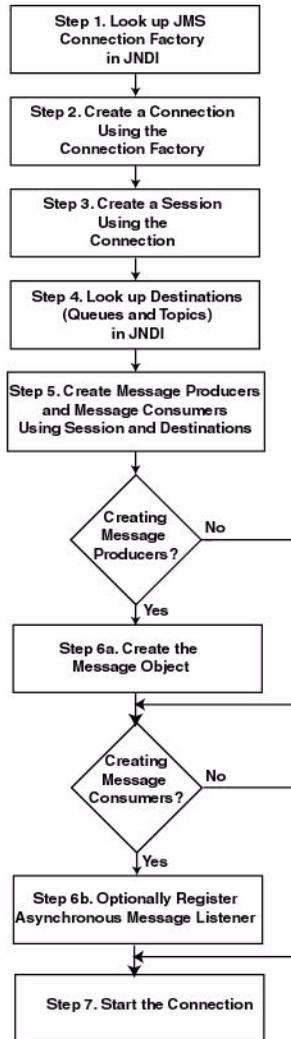
If you want to utilize any of the WebLogic JMS extension classes described in the previous table, also include the following statement on your import list:

```
import weblogic.jms.extensions.*;
```

# Setting Up a JMS Application

Before you can send and receive messages, you must set up a JMS application. The following figure illustrates the steps required to set up a JMS application.

Figure 4-2 Setting Up a JMS Application



The setup steps are described in the following sections. Detailed examples of setting up a Point-to-Point (PTP) and Publish/Subscribe (Pub/Sub) application are also provided. The examples are excerpted from the `examples.jms` package provided with WebLogic Server in the `WL_HOME\samples\server\src\examples\jms` directory, where `WL_HOME` is the top-level directory of your WebLogic Platform installation.

Before proceeding, ensure that the system administrator responsible for configuring WebLogic Server has configured the required JMS features, including the connection factories, JMS servers, and destinations. For more information, see “[Managing JMS](#)” in the *Administration Guide*.

For more information about the JMS classes and methods described in these sections, see “[WebLogic JMS Classes](#)” on page 2-5, or the `javax.jms`, or the `weblogic.jms.ServerSessionPoolFactory`, or the `weblogic.jms.extensions` Javadoc.

For information about setting up transacted applications and JTA user transactions, see “[Using Transactions with WebLogic JMS](#)” on page 5-1.

### Step 1: Look Up a Connection Factory in JNDI

Before you can look up a connection factory, it must be defined as part of the configuration information. WebLogic JMS provides one default connection factory, that is included as part of the configuration by default. The WebLogic JMS system administrator may add or update connection factories during configuration. For information on configuring connection factories and the defaults that are available, see “[Managing JMS](#)” in the *Administration Guide*.

Once the connection factory has been defined, you can look it up by first establishing a JNDI context (`context`) using the `NamingManager.InitialContext()` method. For any application other than a servlet application, you must pass an environment used to create the initial context. For more information, see the `NamingManager.InitialContext()` Javadoc.

Once the context is defined, to look up a connection factory in JNDI, execute one of the following commands, for PTP or Pub/Sub messaging, respectively:

```
QueueConnectionFactory queueConnectionFactory =  
    (QueueConnectionFactory) context.lookup(CF_name);  
  
TopicConnectionFactory topicConnectionFactory =  
    (TopicConnectionFactory) context.lookup(CF_name);
```

The `CF_name` argument specifies the connection factory name defined during configuration.

For more information about the `ConnectionFactory` class, see “[ConnectionFactory](#)” on page 2-6 or the `javax.jms.ConnectionFactory` Javadoc.

## Step 2: Create a Connection Using the Connection Factory

You can create a connection for accessing a queue or topic using the `ConnectionFactory` methods described in the following sections.

For more information about the `Connection` class, see [“Connection” on page 2-8](#) or the [`javax.jms.Connection` Javadoc](#).

### Create a Queue Connection

The `QueueConnectionFactory` provides the following two methods for creating a queue connection:

```
public QueueConnection createQueueConnection(  
    ) throws JMSEException  
  
public QueueConnection createQueueConnection(  
    String userName,  
    String password  
    ) throws JMSEException
```

The first method creates a `QueueConnection`; the second method creates a `QueueConnection` using a specified user identity. In each case, a connection is created in stopped mode and must be started in order to accept messages, as described in [“Step 7: Start the Connection” on page 4-16](#).

For more information about the `QueueConnectionFactory` class methods, see the [`javax.jms.QueueConnectionFactory` Javadoc](#). For more information about the `QueueConnection` class, see the [`javax.jms.QueueConnection` Javadoc](#).

### Create a Topic Connection

The `TopicConnectionFactory` provides the following two methods for creating a topic connection:

```
public TopicConnection createTopicConnection(  
    ) throws JMSEException  
  
public TopicConnection createTopicConnection(  
    String userName,  
    String password  
    ) throws JMSEException
```

The first method creates a `TopicConnection`; the second method creates a `TopicConnection` using a specified user identity. In each case, a connection is created in stopped mode and must be started in order to accept messages, as described in “[Step 7: Start the Connection](#)” on page 4-16.

For more information about the `TopicConnectionFactory` class methods, see the [javax.jms.TopicConnectionFactory](#) Javadoc. For more information about the `TopicConnection` class, see the [javax.jms.TopicConnection](#) Javadoc.

### Step 3: Create a Session Using the Connection

You can create one or more sessions for accessing a queue or topic using the `Connection` methods described in the following sections.

**Note:** A session and its message producers and consumers can only be accessed by one thread at a time. Their behavior is undefined if multiple threads access them simultaneously.

For more information about the `Session` class, see “[Session](#)” on page 2-9 or the [javax.jms.Session](#) Javadoc.

#### Create a Queue Session

The `QueueConnection` class defines the following method for creating a queue session:

```
public QueueSession createQueueSession(  
    boolean transacted,  
    int acknowledgeMode  
) throws JMSException
```

You must specify a boolean argument indicating whether the session will be transacted (`true`) or non-transacted (`false`), and an integer that indicates the acknowledge mode for non-transacted sessions, as described in [Table 2-5, “Acknowledge Modes Used for Non-Transacted Sessions,”](#) on page 2-10. The `acknowledgeMode` attribute is ignored for transacted sessions. In this case, messages are acknowledged when the transaction is committed using the `commit()` method.

For more information about the `QueueConnection` class methods, see the [javax.jms.QueueConnection](#) Javadoc. For more information about the `QueueSession` class, see the [javax.jms.QueueSession](#) Javadoc.

## Create a Topic Session

The `TopicConnection` class defines the following method for creating a topic session:

```
public TopicSession createTopicSession(  
    boolean transacted,  
    int acknowledgeMode  
) throws JMSException
```

You must specify a boolean argument indicating whether the session will be transacted (`true`) or non-transacted (`false`), and an integer that indicates the acknowledge mode for non-transacted sessions, as described in [“Acknowledge Modes Used for Non-Transacted Sessions” on page 2-10](#). The `acknowledgeMode` attribute is ignored for transacted sessions. In this case, messages are acknowledged when the transaction is committed using the `commit()` method.

For more information about the `TopicConnection` class methods, see the [`javax.jms.TopicConnection` Javadoc](#). For more information about the `TopicSession` class, see the [`javax.jms.TopicSession` Javadoc](#).

## Step 4: Look Up a Destination (Queue or Topic)

Before you can look up a destination, the destination must be configured by the WebLogic JMS system administrator, as described in [“Managing JMS” in the \*Administration Guide\*](#).

Once the destination has been configured, you can look up a destination by establishing a JNDI context (`context`), which has already been accomplished in [“Step 1: Look Up a Connection Factory in JNDI” on page 4-6](#), and executing one of the following commands, for PTP or Pub/Sub messaging, respectively:

```
Queue queue = (Queue) context.lookup(Dest_name);  
Topic topic = (Topic) context.lookup(Dest_name);
```

The `Dest_name` argument specifies the destination’s JNDI name defined during configuration.

If you do not use a JNDI namespace, you can use the following `QueueSession` or `TopicSession` method to reference a queue or topic, respectively:

```
public Queue createQueue(  
    String queueName  
    ) throws JMSEException  
  
public Topic createTopic(  
    String topicName  
    ) throws JMSEException
```

The syntax for the `queueName` and/or `topicName` string is `JMS_Server_Name/Destination_Name` (for example, `myjmsserver/mydestination`). To view source code that uses this syntax, refer to the `findqueue()` example in “[Creating Destinations Dynamically](#)” on page 4-49.

**Note:** The `createQueue()` and `createTopic()` methods *do not create* destinations dynamically; they create only references to destinations that already exist. For information about creating destinations dynamically, see “[Creating Destinations Dynamically](#)” on page 4-49.

For more information about these methods, see the `javax.jms.QueueSession` and `javax.jms.TopicSession` Javadoc, respectively.

Once the destination has been defined, you can use the following `Queue` or `Topic` method to access the queue or topic name, respectively:

```
public String getQueueName(  
    ) throws JMSEException  
  
public String getTopicName(  
    ) throws JMSEException
```

To ensure that the queue and topic names are returned in printable format, use the `toString()` method.

For more information about the `Destination` class, see “[Destination](#)” on page 2-13 or the `javax.jms.Destination` Javadoc.

### Server Affinity When Looking Up Destinations

The `createTopic()` and `createQueue()` methods also allow a “`./Destination_Name`” syntax to indicate server affinity when looking up destinations. This will locate destinations that are locally deployed in the same JVM as the JMS connection’s connection factory host. If the name is not on the local JVM an exception is thrown, even though the same name might be deployed on a different JVM.

An application might use this convention to avoid hard-coding the server name when using the `createTopic()` and `createQueue()` methods so that the code can be reused on different JMS servers without requiring any changes.

## Step 5: Create Message Producers and Message Consumers Using the Session and Destinations

You can create message producers and message consumers by passing the destination reference to the `Session` methods described in the following sections.

**Note:** Each consumer receives its own local copy of a message. Once received, you can modify the header field values; however, the message properties and message body are read only. (Attempting to modify the message properties or body at this point will generate a `MessageNotWriteableException`.) You can modify the message body by executing the corresponding message type's `clearbody()` method to clear the existing contents and enable write permission.

For more information about the `MessageProducer` and `MessageConsumer` classes, see “[MessageProducer and MessageConsumer](#)” on page 2-14, or the [javax.jms.MessageProducer](#) and [javax.jms.MessageConsumer](#) Javadocs, respectively.

### Create QueueSenders and QueueReceivers

The `QueueSession` object defines the following methods for creating queue senders and receivers:

```
public QueueSender createSender(
    Queue queue
) throws JMSException

public QueueReceiver createReceiver(
    Queue queue
) throws JMSException

public QueueReceiver createReceiver(
    Queue queue,
    String messageSelector
) throws JMSException
```

You must specify the queue object for the queue sender or receiver being created. You may also specify a message selector for filtering messages. Message selectors are described in more detail in “[Filtering Messages](#)” on page 4-66.

If you pass a value of null to the `createSender()` method, you create an *anonymous producer*. In this case, you must specify the queue name when sending messages, as described in “[Sending Messages](#)” on page 4-22.

Once the queue sender or receiver has been created, you can access the queue name associated with the queue sender or receiver using the following `QueueSender` or `QueueReceiver` method:

```
public Queue getQueue(  
    ) throws JMSEException
```

For more information about the `QueueSession` class methods, see the [javax.jms.QueueSession](#) Javadoc. For more information about the `QueueSender` and `QueueReceiver` classes, see the [javax.jms.QueueSender](#) and [javax.jms.QueueReceiver](#) Javadocs, respectively.

### Create TopicPublishers and TopicSubscribers

The `TopicSession` object defines the following methods for creating topic publishers and topic subscribers:

```
public TopicPublisher createPublisher(  
    Topic topic  
    ) throws JMSEException  
  
public TopicSubscriber createSubscriber(  
    Topic topic  
    ) throws JMSEException  
  
public TopicSubscriber createSubscriber(  
    Topic topic,  
    String messageSelector,  
    boolean noLocal  
    ) throws JMSEException
```

**Note:** The methods described in this section create non-durable subscribers. Non-durable topic subscribers only receive messages sent while they are active. For information about the methods used to create durable subscriptions enabling messages to be retained until all messages are delivered to a durable

subscriber, see [“Setting Up Durable Subscriptions” on page 4-53](#). In this case, durable subscribers only receive messages that are published after the subscriber has subscribed.

You must specify the topic object for the publisher or subscriber being created. You may also specify a message selector for filtering messages and a `noLocal` flag (described later in this section). Message selectors are described in more detail in [“Filtering Messages” on page 4-66](#).

If you pass a value of null to the `createPublisher()` method, you create an *anonymous producer*. In this case, you must specify the topic name when sending messages, as described in [“Sending Messages” on page 4-22](#).

An application can have a JMS connection that it uses to both publish and subscribe to the same topic. Because topic messages are delivered to all subscribers, the application can receive messages it has published itself. To prevent this behavior, a JMS application can set a `noLocal` flag to `true`.

Once the topic publisher or subscriber has been created, you can access the topic name associated with the topic publisher or subscriber using the following `TopicPublisher` or `TopicSubscriber` method:

```
Topic getTopic(  
    ) throws JMSEException
```

In addition, you can access the `noLocal` variable setting associated with the topic subscriber using the following `TopicSubscriber` method:

```
boolean getNoLocal(  
    ) throws JMSEException
```

For more information about the `TopicSession` class methods, see the [`javax.jms.TopicSession` Javadoc](#). For more information about the `TopicPublisher` and `TopicSubscriber` classes, see the [`javax.jms.TopicPublisher` and `javax.jms.TopicSubscriber` Javadocs](#), respectively.

## Step 6a: Create the Message Object (Message Producers)

**Note:** This step applies to message producers only.

To create the message object, use one of the following `Session` or `WLSession` class methods:

### ■ Session Methods

**Note:** These methods are inherited by both the `QueueSession` and `TopicSession` subclasses.

```
public BytesMessage createBytesMessage(  
    ) throws JMSEException  
  
public MapMessage createMapMessage(  
    ) throws JMSEException  
  
public Message createMessage(  
    ) throws JMSEException  
  
public ObjectMessage createObjectMessage(  
    ) throws JMSEException  
  
public ObjectMessage createObjectMessage(  
    Serializable object  
    ) throws JMSEException  
  
public StreamMessage createStreamMessage(  
    ) throws JMSEException  
  
public TextMessage createTextMessage(  
    ) throws JMSEException  
  
public TextMessage createTextMessage(  
    String text  
    ) throws JMSEException
```

### ■ WLSession Method

```
public XMLMessage createXMLMessage(  
    String text  
    ) throws JMSEException
```

For more information about the `Session` and `WLSession` class methods, see the [javax.jms.Session](#) and [weblogic.jms.extensions.WLSession](#) Javadocs, respectively. For more information about the `Message` class and its methods, see “Message” on page 2-16, or the [javax.jms.Message](#) Javadoc.

## Step 6b: Optionally Register an Asynchronous Message Listener (Message Consumers)

**Note:** This step applies to message consumers only.

To receive messages asynchronously, you must register an asynchronous message listener by performing the following steps:

1. Implement the `javax.jms.MessageListener` interface, which includes an `onMessage()` method.

**Note:** For an example of the `onMessage()` method interface, see “[Example: Setting Up a PTP Application](#)” on page 4-16.

If you wish to issue the `close()` method within an `onMessage()` method call, the system administrator must select the Allow Close In OnMessage check box when configuring the connection factory. For more information on configuring JMS, see “[Managing JMS](#)” in the *Administration Guide*.

2. Set the message listener using the following `MessageConsumer` method, passing the listener information as an argument:

```
public void setMessageListener(  
    MessageListener listener  
    ) throws JMSEException
```

3. Optionally, implement an exception listener on the session to catch exceptions, as described in “[Defining a Session Exception Listener](#)” on page 4-48.

You can unset a message listener by calling the `MessageListener()` method with a value of null.

Once a message listener has been defined, you can access it by calling the following `MessageConsumer` method:

```
public MessageListener getMessageListener(  
    ) throws JMSEException
```

**Note:** WebLogic JMS guarantees that multiple `onMessage()` calls for the same session will not be executed simultaneously.

If a message consumer is closed by an administrator or as the result of a server failure, a `ConsumerClosedException` is delivered to the session exception listener, if one has been defined. In this way, a new message consumer can be created, if necessary. For information about defining a session exception listener, see [“Defining a Session Exception Listener” on page 4-48](#).

The `MessageConsumer` class methods are inherited by the `QueueReceiver` and `TopicSubscriber` classes. For additional information about the `MessageConsumer` class methods, see [“MessageProducer and MessageConsumer” on page 2-14](#) or the `javax.jms.MessageConsumer` Javadoc.

## Step 7: Start the Connection

You start the connection using the `Connection` class `start()` method.

For additional information about starting, stopping, and closing a connection, see [“Starting, Stopping, and Closing a Connection” on page 4-46](#) or the `javax.jms.Connection` Javadoc.

## Example: Setting Up a PTP Application

The following example is excerpted from the `examples.jms.queue.QueueSend` example, provided with WebLogic Server in the `WL_HOME\samples\server\src\examples\jms\queue` directory, where `WL_HOME` is the top-level directory of your WebLogic Platform installation. The `init()` method shows how to set up and start a `QueueSession` for a JMS application. The following shows the `init()` method, with comments describing each setup step.

Define the required variables, including the JNDI context, JMS connection factory, and queue static variables.

```
public final static String JNDI_FACTORY=
    "weblogic.jndi.WLInitialContextFactory";
public final static String JMS_FACTORY=
    "weblogic.examples.jms.QueueConnectionFactory";
public final static String
    QUEUE="weblogic.examples.jms.exampleQueue";

private QueueConnectionFactory qconFactory;
```

```
private QueueConnection qcon;
private QueueSession qsession;
private QueueSender qsender;
private Queue queue;
private TextMessage msg;
```

Set up the JNDI initial context, as follows:

```
InitialContext ic = getInitialContext(args[0]);
    .
    .
    .
private static InitialContext getInitialContext(
    String url
) throws NamingException
{
    Hashtable env = new Hashtable();
    env.put(Context.INITIAL_CONTEXT_FACTORY, JNDI_FACTORY);
    env.put(Context.PROVIDER_URL, url);
    return new InitialContext(env);
}
```

**Note:** When setting up the JNDI initial context for an EJB or servlet, use the following method:

```
Context ctx = new InitialContext();
```

Create all the necessary objects for sending messages to a JMS queue. The `ctx` object is the JNDI initial context passed in by the `main()` method.

```
public void init(
    Context ctx,
    String queueName
) throws NamingException, JMSEException
{
```

**Step 1** Look up a connection factory in JNDI.

```
qconFactory = (QueueConnectionFactory) ctx.lookup(JMS_FACTORY);
```

**Step 2** Create a connection using the connection factory.

```
qcon = qconFactory.createQueueConnection();
```

**Step 3** Create a session using the connection. The following code defines the session as non-transacted and specifies that messages will be acknowledged automatically. For more information about transacted sessions and acknowledge modes, see [“Session” on page 2-9](#).

## 4 Developing a WebLogic JMS Application

---

```
qsession = qcon.createQueueSession(false,
    Session.AUTO_ACKNOWLEDGE);
```

Step 4 Look up a destination (queue) in JNDI.

```
queue = (Queue) ctx.lookup(queueName);
```

Step 5 Create a reference to a message producer (queue sender) using the session and destination (queue).

```
qsender = qsession.createSender(queue);
```

Step 6 Create the message object.

```
msg = qsession.createTextMessage();
```

Step 7 Start the connection.

```
qcon.start();
}
```

The `init()` method for the `examples.jms.queue.QueueReceive` example is similar to the `QueueSend` `init()` method shown previously, with the one exception. Steps 5 and 6 would be replaced by the following code, respectively:

```
greceiver = qsession.createReceiver(queue);
greceiver.setMessageListener(this);
```

In the first line, instead of calling the `createSender()` method to create a reference to the queue sender, the application calls the `createReceiver()` method to create the queue receiver.

In the second line, the message consumer registers an asynchronous message listener.

When a message is delivered to the queue session, it is passed to the `examples.jms.QueueReceive.onMessage()` method. The following code excerpt shows the `onMessage()` interface from the `QueueReceive` example:

```
public void onMessage(Message msg)
{
    try {
        String msgText;
        if (msg instanceof TextMessage) {
            msgText = ((TextMessage)msg).getText();
        } else { // If it is not a TextMessage...
            msgText = msg.toString();
        }

        System.out.println("Message Received: "+ msgText );
    }
}
```

```

    if (msgText.equalsIgnoreCase("quit")) {
        synchronized(this) {
            quit = true;
            this.notifyAll(); // Notify main thread to quit
        }
    }
} catch (JMSEException jmse) {
    jmse.printStackTrace();
}
}

```

The `onMessage()` method processes messages received through the queue receiver. The method verifies that the message is a `TextMessage` and, if it is, prints the text of the message. If `onMessage()` receives a different message type, it uses the message's `toString()` method to display the message contents.

**Note:** It is good practice to verify that the received message is the type expected by the handler method.

For more information about the JMS classes used in this example, see [“WebLogic JMS Classes” on page 2-5](#) or the `javax.jms` Javadoc.

## Example: Setting Up a Pub/Sub Application

The following example is excerpted from the `examples.jms.topic.TopicSend` example, provided with WebLogic Server in the `WL_HOME\samples\server\src\examples\jms\topic` directory, where `WL_HOME` is the top-level directory of your WebLogic Platform installation. The `init()` method shows how to set up and start a topic session for a JMS application. The following shows the `init()` method, with comments describing each setup step.

Define the required variables, including the JNDI context, JMS connection factory, and topic static variables.

```

public final static String JNDI_FACTORY=
    "weblogic.jndi.WLInitialContextFactory";
public final static String JMS_FACTORY=
    "weblogic.examples.jms.TopicConnectionFactory";
public final static String
    TOPIC="weblogic.examples.jms.exampleTopic";

protected TopicConnectionFactory tconFactory;
protected TopicConnection tcon;

```

```
protected TopicSession tsession;
protected TopicPublisher tpublisher;
protected Topic topic;
protected TextMessage msg;
```

Set up the JNDI initial context, as follows:

```
InitialContext ic = getInitialContext(args[0]);
    .
    .
    .
private static InitialContext getInitialContext(
    String url
) throws NamingException
{
    Hashtable env = new Hashtable();
    env.put(Context.INITIAL_CONTEXT_FACTORY, JNDI_FACTORY);
    env.put(Context.PROVIDER_URL, url);
    return new InitialContext(env);
}
```

**Note:** When setting up the JNDI initial context for a servlet, use the following method:

```
Context ctx = new InitialContext();
```

Create all the necessary objects for sending messages to a JMS queue. The `ctx` object is the JNDI initial context passed in by the `main()` method.

```
public void init(
    Context ctx,
    String topicName
) throws NamingException, JMSEException
{
```

**Step 1** Look up a connection factory using JNDI.

```
tconFactory =
    (TopicConnectionFactory) ctx.lookup(JMS_FACTORY);
```

**Step 2** Create a connection using the connection factory.

```
tcon = tconFactory.createTopicConnection();
```

**Step 3** Create a session using the connection. The following defines the session as non-transacted and specifies that messages will be acknowledged automatically. For more information about setting session transaction and acknowledge modes, see [“Session” on page 2-9](#).

```
tsession = tcon.createTopicSession(false,
    Session.AUTO_ACKNOWLEDGE);
```

Step 4 Look up the destination (topic) using JNDI.

```
topic = (Topic) ctx.lookup(topicName);
```

Step 5 Create a reference to a message producer (topic publisher) using the session and destination (topic).

```
tpublisher = tsession.createPublisher(topic);
```

Step 6 Create the message object.

```
msg = tsession.createTextMessage();
```

Step 7 Start the connection.

```
tcon.start();
}
```

The `init()` method for the `examples.jms.topic.TopicReceive` example is similar to the `TopicSend` `init()` method shown previously with one exception. Steps 5 and 6 would be replaced by the following code, respectively:

```
tsubscriber = tsession.createSubscriber(topic);
tsubscriber.setMessageListener(this);
```

In the first line, instead of calling the `createPublisher()` method to create a reference to the topic publisher, the application calls the `createSubscriber()` method to create the topic subscriber.

In the second line, the message consumer registers an asynchronous message listener.

When a message is delivered to the topic session, it is passed to the `examples.jms.TopicSubscribe.onMessage()` method. The `onMessage()` interface for the `TopicReceive` example is the same as the `QueueReceive.onMessage()` interface, as described in [“Example: Setting Up a PTP Application” on page 4-16](#).

For more information about the JMS classes used in this example, see [“WebLogic JMS Classes” on page 2-5](#) or the `javax.jms` Javadoc.

# Sending Messages

Once you have set up the JMS application as described in [“Setting Up a JMS Application” on page 4-4](#), you can send messages. To send a message, you must perform the following steps:

1. Create a message object.
2. Define a message.
3. Send the message to a destination.

For more information about the JMS classes for sending messages and the message types, see the `javax.jms.Message` Javadoc. For information about receiving messages, see [“Receiving Messages” on page 4-29](#).

## Step 1: Create a Message Object

This step has already been accomplished as part of the client setup procedure, as described in [“Step 6a: Create the Message Object \(Message Producers\)” on page 4-13](#).

## Step 2: Define a Message

This step *may* have been accomplished when setting up an application, as described in [“Step 6a: Create the Message Object \(Message Producers\)” on page 4-13](#). Whether or not this step has already been accomplished depends on the method that was called to create the message object. For example, for `TextMessage` and `ObjectMessage` types, when you create a message object, you have the option of defining the message when you create the message object.

If a value has been specified and you do not wish to change it, you can proceed to step 3.

If a value has not been specified or if you wish to change an existing value, you can define a value using the appropriate `set` method. For example, the method for defining the text of a `TextMessage` is as follows:

```
public void setText(  
    String string  
) throws JMSEException
```

**Note:** Messages can be defined as null.

Subsequently, you can clear the message body using the following method:

```
public void clearBody(  
) throws JMSEException
```

For more information about the methods used to define messages, see the [javax.jms.Session](#) Javadoc.

## Step 3: Send the Message to a Destination

You can send a message to a destination using a message producer—queue sender (PTP) or topic publisher (Pub/Sub)—and the methods described in the following sections. The `Destination` and `MessageProducer` objects were created when you set up the application, as described in “[Setting Up a JMS Application](#)” on page 4-4.

**Note:** If multiple topic subscribers are defined for the same topic, each subscriber will receive its own local copy of a message. Once received, you can modify the header field values; however, the message properties and message body are read only. You can modify the message body by executing the corresponding message type’s `clearbody()` method to clear the existing contents and enable write permission.

For more information about the `MessageProducer` class, see “[MessageProducer and MessageConsumer](#)” on page 2-14 or the [javax.jms.MessageProducer](#) Javadoc.

### Send a Message Using Queue Sender

You can send messages using the following `QueueSender` methods:

```
public void send(  
    Message message  
) throws JMSEException  
  
public void send(  
    Message message,  
    int deliveryMode,
```

```
        int priority,
        long timeToLive
    ) throws JMSEException

    public void send(
        Queue queue,
        Message message
    ) throws JMSEException

    public void send(
        Queue queue,
        Message message,
        int deliveryMode,
        int priority,
        long timeToLive
    ) throws JMSEException
```

You must specify a message. You may also specify the queue name (for anonymous message producers), delivery mode (`DeliveryMode.PERSISTENT` or `DeliveryMode.NON_PERSISTENT`), priority (0-9), and time-to-live (in milliseconds). If not specified, the delivery mode, priority, and time-to-live attributes are set to one of the following:

- Connection factory or destination override configuration attributes defined for the producer, as described “[Managing JMS](#)” in the *Administration Guide*.
- Values specified using the message producer’s set methods, as described in “[Setting Message Producer Attributes](#)” on page 4-26.

**Note:** WebLogic JMS also provides the following proprietary attributes, as described in “[Setting Message Producer Attributes](#)” on page 4-26.

- `TimeToDeliver` (that is, birth time), which represents the delay before a sent message is made visible on its target destination.
- `RedeliveryLimit`, which determines the number of times a message is redelivered after a recover or rollback.

If you define the delivery mode as `PERSISTENT`, you should configure a backing store for the destination, as described in “[Managing JMS](#)” in the *Administration Guide*.

**Note:** If no backing store is configured, then the delivery mode is changed to `NON_PERSISTENT` and messages are not written to the persistent store.

If the queue sender is an anonymous producer (that is, if when the queue was created, the name was set to null), then you must specify the queue name (using one of the last two methods) to indicate where to deliver messages. For more information about defining anonymous producers, see [“Create QueueSenders and QueueReceivers” on page 4-11](#).

For example, the following code sends a persistent message with a priority of 4 and a time-to-live of one hour:

```
QueueSender.send(message, DeliveryMode.PERSISTENT, 4, 3600000);
```

For additional information about the `QueueSender` class methods, see the [javax.jms.QueueSender](#) Javadoc.

## Send a Message Using TopicPublisher

You can send messages using the following `TopicPublisher` methods:

```
public void publish(
    Message message
) throws JMSEException

public void publish(
    Message message,
    int deliveryMode,
    int priority,
    long timeToLive
) throws JMSEException

public void publish(
    Topic topic,
    Message message
) throws JMSEException

public void publish(
    Topic topic,
    Message message,
    int deliveryMode,
    int priority,
    long timeToLive
) throws JMSEException
```

You must provide a message. You may also specify the topic name, delivery mode (`DeliveryMode.PERSISTENT` or `DeliveryMode.NON_PERSISTENT`), priority (0-9), and time-to-live (in milliseconds). If not specified, the delivery mode, priority, and time-to-live attributes are set to one of the following:

- Connection factory or destination override configuration attributes defined for the producer, as described “[Managing JMS](#)” in the *Administration Guide*.
- Values specified using the message producer’s set methods, as described in “[Setting Message Producer Attributes](#)” on page 4-26.

**Note:** WebLogic JMS also provides the following proprietary attributes, as described in “[Setting Message Producer Attributes](#)” on page 4-26.

- `TimeToDeliver` (that is, birth time), which represents the delay before a sent message is made visible on its target destination.
- `RedeliveryLimit`, which determines the number of times a message is redelivered after a recover or rollback.

If you define the delivery mode as `PERSISTENT`, you should configure a backing store, as described in “[Managing JMS](#)” in the *Administration Guide*.

**Note:** If no backing store is configured, then the delivery mode is changed to `NON_PERSISTENT` and no messages are stored.

If the topic publisher is an anonymous producer (that is, if when the topic was created, the name was set to null), then you must specify the topic name (using either of the last two methods) to indicate where to deliver messages. For more information about defining anonymous producers, see “[Create TopicPublishers and TopicSubscribers](#)” on page 4-12.

For example, the following code sends a persistent message with a priority of 4 and a time-to-live of one hour:

```
TopicPublisher.publish(message, DeliveryMode.PERSISTENT,  
    4, 3600000);
```

For more information about the `TopicPublisher` class methods, see the [javax.jms.TopicPublisher](#) Javadoc.

## Setting Message Producer Attributes

As described in the previous section, when sending a message, you can optionally specify the delivery mode, priority, and time-to-live values. If not specified, these attributes are set to the connection factory configuration attributes, as described “[Managing JMS](#)” in the *Administration Guide*.

Alternatively, you can set the delivery mode, priority, time-to-deliver, time-to-live, and redelivery delay (timeout), and redelivery limit values dynamically using the message producer's set methods. The following table lists the message producer set and get methods for each dynamically configurable attribute.

**Note:** The delivery mode, priority, time-to-live, time-to-deliver, redelivery delay (timeout), and redelivery limit attribute settings can be overridden by the destination using the Delivery Mode Override, Priority Override, Time To Live Override, Time To Deliver Override, Redelivery Delay Override, and Redelivery Limit configuration attributes, as described in [Administration Console Online Help](#).

**Table 4-2 Message Producer Set and Get Methods**

Attribute	Set Method	Get Method
Delivery Mode	<code>public void setDeliveryMode(int deliveryMode) throws JMSEException</code>	<code>public int getDeliveryMode() throws JMSEException</code>
Priority	<code>public void setPriority(int defaultPriority) throws JMSEException</code>	<code>public int getPriority() throws JMSEException</code>
Time-To-Live	<code>public void setTimeToLive(long timeToLive) throws JMSEException</code>	<code>public long getTimeToLive() throws JMSEException</code>
Time-To-Deliver	<code>public void setTimeToDeliver(long timeToDeliver) throws JMSEException</code>	<code>public long getTimeToDeliver() throws JMSEException</code>
Redelivery Limit	<code>public void setRedeliveryLimit(int redeliveryLimit) throws JMSEException</code>	<code>public int getRedeliveryLimit() throws JMSEException</code>

**Note:** JMS defines optional `MessageProducer` methods for disabling the message ID and timestamp information. However, these methods are ignored by WebLogic JMS.

For more information about the `MessageProducer` class methods, see Sun's [javax.jms.MessageProducer Javadoc](#) or the [weblogic.jms.extensions.WLMessageProducer Javadoc](#).

### Example: Sending Messages Within a PTP Application

The following example is excerpted from the `examples.jms.queue.QueueSend` example, provided with WebLogic Server in the `WL_HOME\samples\server\src\examples\jms\queue` directory, where `WL_HOME` is the top-level directory of your WebLogic Platform installation. The example shows the code required to create a `TextMessage`, set the text of the message, and send the message to a queue.

```
msg = qsession.createTextMessage();
    .
    .
    .
public void send(
    String message
) throws JMSEException
{
    msg.setText(message);
    qsender.send(msg);
}
```

For more information about the `QueueSender` class and methods, see the [javax.jms.QueueSender](#) Javadoc.

### Example: Sending Messages Within a Pub/Sub Application

The following example is excerpted from the `examples.jms.topic.TopicSend` example, provided with WebLogic Server in the `WL_HOME\samples\server\src\examples\jms\topic` directory, where `WL_HOME` is the top-level directory of your WebLogic Platform installation. It shows the code required to create a `TextMessage`, set the text of the message, and send the message to a topic.

```
msg = tsession.createTextMessage();
    .
    .
    .
public void send(
    String message
) throws JMSEException
```

```
{  
    msg.setText(message);  
    tpublisher.publish(msg);  
}
```

For more information about the `TopicPublisher` class and methods, see the [javax.jms.TopicPublisher Javadoc](#).

## Receiving Messages

Once you have set up the JMS application as described in “[Setting Up a JMS Application](#)” on page 4-4, you can receive messages.

To receive a message, you must create the receiver object and specify whether you want to receive messages asynchronously or synchronously, as described in the following sections.

The order in which messages are received can be controlled by the following:

- Message delivery attributes (delivery mode and sorting criteria) defined during configuration, as described in “[Managing JMS](#)” in the *Administration Guide*, or as part of the `send()` method, as described in “[Sending Messages](#)” on page 4-22.
- Destination sort order set using destination keys, as described in “[Managing JMS](#)” in the *Administration Guide*.

Once received, you can modify the header field values; however, the message properties and message body are read-only. You can modify the message body by executing the corresponding message type’s `clearbody()` method to clear the existing contents and enable write permission.

For more information about the JMS classes for receiving messages and the message types, see the [javax.jms.Message Javadoc](#). For information about sending messages, see “[Sending Messages](#)” on page 4-22.

# Receiving Messages Asynchronously

This procedure is described within the context of setting up the application. For more information, see [“Step 6b: Optionally Register an Asynchronous Message Listener \(Message Consumers\)”](#) on page 4-15.

You can control the maximum number of messages that may exist for an asynchronous session and that have not yet been passed to the message listener by setting the Messages Maximum attribute when configuring the JMS connection factory.

## Asynchronous Message Pipeline

If messages are produced faster than asynchronous message listeners (consumers) can consume them, a JMS server will push multiple unconsumed messages in a batch to another session with available asynchronous message listeners. These in-flight messages are sometimes referred to as the *message pipeline*, or in some JMS vendors as the *message backlog*. The pipeline or backlog size is the number of messages that have accumulated on an asynchronous consumer, but which have not been passed to a message listener.

### Configuring a Message Pipeline

You can control a client’s maximum pipeline size by configuring the Messages Maximum attribute on the client’s connection factory, which is defined as the “maximum number of messages that can exist for an asynchronous consumer and that have not yet been passed to the message listener”. The default setting is *10*.

For more information on configuring a JMS connection factory, see [“Managing JMS”](#) in the *Administration Guide*.

### Behavior of Pipelined Messages

Once a message pipeline is configured, it will exhibit the following behavior:

- **Statistics** — JMS monitoring statistics reports backlogged messages in a message pipeline as pending (for queues and durable subscribers) until they are either committed or acknowledged.
- **Performance** — Increasing the Messages Maximum pipeline size may improve performance for high-throughput applications. Note that a larger pipeline will increase client memory usage, as the pending pipelined messages accumulate on the client JVM before the asynchronous consumer’s listener is called.

- **Sorting** — Messages in an asynchronous consumer's pipeline are not sorted according to the consumer destination's configured sort order; instead, they remain in the order in which they are pushed from the JMS server. For example, if a destination is configured to sort by priority, high priority messages will not jump ahead of low priority messages that have already been pushed into an asynchronous consumer's pipeline.

**Notes:** The Messages Maximum pipeline size setting on the connection factory is not related to the Messages Maximum quota settings on JMS servers and destinations.

Pipelined messages are sometimes aggregated into a single message on the network transport. If the messages are sufficiently large, the aggregate size of the data written may exceed the maximum value for the transport, which may cause undesirable behavior. For example, the `tcp` protocol sets a default maximum message size of 10,000,000 bytes, and is configurable on the server with the `MaxT3MessageSize` attribute. This means that if ten 2 megabyte messages are pipelined, the `tcp` limit may be exceeded.

## Receiving Messages Synchronously

To receive messages synchronously, use the following `MessageConsumer` methods:

```
public Message receive(  
    ) throws JMSEException  
  
public Message receive(  
    long timeout  
    ) throws JMSEException  
  
public Message receiveNoWait(  
    ) throws JMSEException
```

In each case, the application receives the next message produced. If you call the `receive()` method with no arguments, the call blocks indefinitely until a message is produced or the application is closed. Alternatively, you can pass a timeout value to specify how long to wait for a message. If you call the `receive()` method with a value of 0, the call blocks indefinitely. The `receiveNoWait()` method receives the next message if one is available, or returns null; in this case, the call does not block.

The `MessageConsumer` class methods are inherited by the `QueueReceiver` and `TopicSubscriber` classes. For additional information about the `MessageConsumer` class methods, see the [javax.jms.MessageConsumer](#) Javadoc.

### Example: Receiving Messages Synchronously Within a PTP Application

The following example is excerpted from the `examples.jms.queue.QueueReceive` example, provided with WebLogic Server in the `WL_HOME\samples\server\src\examples\jms\queue` directory, where `WL_HOME` is the top-level directory of your WebLogic Platform installation. Rather than set a message listener, you would call `qreceiver.receive()` for each message. For example:

```
qreceiver = qsession.createReceiver(queue);
qreceiver.receive();
```

The first line creates the queue receiver on the queue. The second line executes a `receive()` method. The `receive()` method blocks and waits for a message.

### Example: Receiving Messages Synchronously Within a Pub/Sub Application

The following example is excerpted from the `examples.jms.topic.TopicReceive` example, provided with WebLogic Server in the `WL_HOME\samples\server\src\examples\jms\topic` directory, where `WL_HOME` is the top-level directory of your WebLogic Platform installation. Rather than set a message listener, you would call `tsubscriber.receive()` for each message.

For example:

```
tsubscriber = tsession.createSubscriber(topic);
Message msg = tsubscriber.receive();
msg.acknowledge();
```

The first line creates the topic subscriber on the topic. The second line executes a `receive()` method. The `receive()` method blocks and waits for a message.

## Recovering Received Messages

**Note:** This section applies only to non-transacted sessions for which the acknowledge mode is set to `CLIENT_ACKNOWLEDGE`, as described in [Table 2-5, “Acknowledge Modes Used for Non-Transacted Sessions,”](#) on page 2-10. Synchronously received `AUTO_ACKNOWLEDGE` messages may not be recovered; they have already been acknowledged.

An application can request that JMS redeliver messages (unacknowledge them) using the following method:

```
public void recover(  
    ) throws JMSEException
```

The `recover()` method performs the following steps:

- Stops message delivery for the session
- Tags all messages that have not been acknowledged (but may have been delivered) as redelivered
- Resumes sending messages starting from the first unacknowledged message for that session

Messages in queues are not necessarily redelivered in the same order that they were originally delivered, nor to the same queue consumers.

## Acknowledging Received Messages

**Note:** This section applies only to non-transacted sessions for which the acknowledge mode is set to `CLIENT_ACKNOWLEDGE`, as described in [Table 2-5, “Acknowledge Modes Used for Non-Transacted Sessions,”](#) on page 2-10.

To acknowledge a received message, use the following `Message` method:

```
public void acknowledge(  
    ) throws JMSEException
```

The `acknowledge()` method depends on how the connection factory’s `Acknowledge Policy` attribute is configured, as follows:

- The default policy of “All” specifies that calling `acknowledge` on a message acknowledges all unacknowledged messages received on the session.
- The “Previous” policy specifies that calling `acknowledge` on a message acknowledges only unacknowledged messages up to, and including, the given message. Messages that are not acknowledged may be redelivered to the client.

This method is effective only when issued by a non-transacted session for which the `acknowledge` mode is set to `CLIENT_ACKNOWLEDGE`. Otherwise, the method is ignored.

# Releasing Object Resources

When you have finished using the connection, session, message producer or consumer, connection consumer, or queue browser created on behalf of a JMS application, you should explicitly close them to release the resources.

Enter the `close()` method to close JMS objects, as follows:

```
public void close(  
    ) throws JMSEException
```

When closing an object:

- The call blocks until the method call completes or until any outstanding asynchronous receiver `onMessage()` calls complete.
- All associated sub-objects are also closed. For example, when closing a session, all associated message producers and consumers are also closed. When closing a connection, all associated sessions are also closed.

For more information about the impact of the `close()` method for each object, see the appropriate `javax.jms` Javadoc. In addition, for more information about the connection or Session `close()` method, see [“Starting, Stopping, and Closing a Connection” on page 4-46](#) or [“Closing a Session” on page 4-49](#), respectively.

The following example is excerpted from the `examples.jms.queue.QueueSend` example, provided with WebLogic Server in the `WL_HOME\samples\server\src\examples\jms\queue` directory, where `WL_HOME` is the top-level directory of your WebLogic Platform installation. This example shows the code required to close the message consumer, session, and connection objects.

```
public void close(  
    ) throws JMSEException  
{  
    qreceiver.close();  
    qsession.close();  
    qcon.close();  
}
```

In the `QueueSend` example, the `close()` method is called at the end of `main()` to close objects and free resources.

## Managing Rolled Back, Recovered, or Expired Messages

The following sections describe how to manage rolled back or recovered messages:

- [Setting a Redelivery Delay for Messages](#)
- [Setting a Redelivery Limit for Messages](#)
- [Passive Message Expiration Policy](#)

### Setting a Redelivery Delay for Messages

You can delay the redelivery of messages when a temporary, external condition prevents an application from properly handling a message. This allows an application to temporarily inhibit the receipt of “poison” messages that it cannot currently handle. When a message is rolled back or recovered, the redelivery delay is the amount of time a message is put aside before an attempt is made to redeliver the message.

If JMS immediately redelivers the message, the error condition may not be resolved and the application may still not be able to handle the message. However, if an application is configured for a redelivery delay, then when it rolls back or recovers a message, the message is set aside until the redelivery delay has passed, at which point the messages are made available for redelivery.

All messages consumed and subsequently rolled back or recovered by a session receive the redelivery delay for that session at the time of rollback or recovery. Messages consumed by multiple sessions as part of a single user transaction will receive different redelivery delays as a function of the session that consumed the individual messages. Messages that are left unacknowledged or uncommitted by a client, either intentionally or as a result of a failure, are not assigned a redelivery delay.

### Setting a Redelivery Delay

A session inherits the redelivery delay from its connection factory when the session is created. The `RedeliveryDelay` attribute of a connection factory is configured using the Administration Console. For more information, see “[JMS Connection Factories](#)” in the Administration Console Online Help.

The application that creates the session can then override the connection factory setting using WebLogic-specific extensions to the `javax.jms.Session` interface. The session attribute is dynamic and can be changed at any time. Changing the session redelivery delay affects all messages consumed and rolled back (or recovered) by that session after the change.

The method for setting the redelivery delay on a session is provided through the `weblogic.jms.extensions.WLSession` interface, which is an extension to the `javax.jms.Session` interface. To define a redelivery delay for a session, use the following methods:

```
public void setRedeliveryDelay(  
    long redeliveryDelay  
) throws JMSEException;  
  
public long getRedeliveryDelay(  
) throws JMSEException;
```

For more information on the `WLSession` class, refer to the [weblogic.jms.extensions.WLSession](#) Javadoc.

### Overriding the Redelivery Delay on a Destination

Regardless of what redelivery delay is set on the session, the destination where a message is being rolled back or recovered can override the setting. The redelivery delay override applied to the redelivery of a message is the one in effect at the time a message is rolled back or recovered.

The `RedeliveryDelayOverride` attribute of a destination is configured using the Administration Console. For more information, see [“JMS Destinations”](#) in the Administration Console Online Help.

## Setting a Redelivery Limit for Messages

You can specify a limit on the number of times that WebLogic JMS will attempt to redeliver a message to an application. Once WebLogic JMS fails to redeliver a message to a destination for a specific number of times, the message can be redirected to an error destination that is associated to the message destination. If the redelivery limit is configured, but no error destination is configured, then persistent or non-persistent messages are simply deleted when they reach their redelivery limit.

Alternatively, you can set the redelivery limit value dynamically using the message producer’s `set` method, as described in [“Setting Message Producer Attributes”](#) on page 4-26.

## Configuring a Message Redelivery Limit

When a destination’s attempts to redeliver a message to a consumer reaches a specified redelivery limit, then the destination deems the message undeliverable. The `RedeliveryLimit` attribute is set on a destination and is configurable using the Administration Console. This setting overrides the redelivery limit set on the message producer. For more information, see [“JMS Destinations”](#) in the Administration Console Online Help.

## Configuring an Error Destination for Undelivered Messages

If an error destination is configured for undelivered messages, then when a message has been deemed undeliverable, the message will be redirected to a specified error destination. The error destination can be either a queue or a topic, and it must be configured on the same JMS server as the destination for which it is defined. If no error destination is configured, then undeliverable messages are simply deleted.

The `ErrorDestination` attribute is configured using the Administration Console. For more information, see [“JMS Destinations”](#) in the Administration Console Online Help.

If a message redelivery attempt has already reached its specified redelivery limit, but its error destination has also reached its maximum quota, then the message is deemed undeliverable and is dropped. Non-persistent messages are deleted, while persistent messages remain in the store and will reappear in their originating destination (not the error destination) when the server is restarted. In either case, a log message is generated. To prevent the log file from becoming clogged, the log message is only generated once per error destination every five minutes, until the error condition is resolved.

## Passive Message Expiration Policy

WebLogic JMS has only a *passive* message expiration policy, in that it does not actively search for or remove expired messages from the system. As such, expired messages are only deleted from the system when:

- The messages are about to be consumed
- The JMS server is rebooted

You should be aware that since no active searches are made for expired messages, they can accumulate on the system and strain system resources.

## Setting Message Delivery Times

You can schedule message deliveries to an application for specific times in the future. Message deliveries can be deferred for short periods of time (such as seconds or minutes) or for long stretches of time (for example, hours later for batch processing). Until that delivery time, the message is essentially invisible until it is delivered, allowing you to schedule work at a particular time in the future.

Messages are not sent on a recurring basis; they are sent only once. In order to send messages on a recurring basis, a received scheduled message must be sent back to its original destination. Typically, the receive, the send, and any associated work should be under the same transaction to ensure exactly-once semantics.

## Setting a Delivery Time on Producers

Support for setting and getting a time-to-deliver on an individual producer is provided through the `weblogic.jms.extensions.WLMessageProducer` interface, which is an extension to the `javax.jms.MessageProducer` interface. To define a time-to-deliver on an individual producer, use the following methods:

```
public void setTimeToDeliver(  
    long timeToDeliver  
) throws JMSEException;  
  
public long getTimeToDeliver(  
) throws JMSEException;
```

For more information on the `WLMessageProducer` class, refer to the [weblogic.jms.extensions.WLMessageProducer](#) Javadoc.

## Setting a Delivery Time on Messages

The `DeliveryTime` is a JMS message header field that defines the earliest absolute time at which the message can be delivered. That is, the message is held by the messaging system and is not given to any consumers until that time.

As a JMS header field, the `DeliveryTime` can be used to sort messages in a destination or to select messages. For purposes of data type conversion, the delivery time is stored as a long integer.

**Note:** Setting a delivery time value on a message has no effect on the `DeliveryTime` field, because JMS will always override the value with the producer's value when the message is sent or published. The message delivery time methods described here are similar to other JMS message fields that are set through the producer, including the delivery mode, priority, time-to-deliver, time-to-live, redelivery delay, and redelivery limit fields. Specifically, the setting of these fields is reserved for JMS providers, including WebLogic JMS.

The support for setting and getting the delivery time on a message is provided through the `weblogic.jms.extensions.WLMessage` interface, which is an extension to the `javax.jms.Message` interface. To define a delivery time on a message, use the following methods:

```
public void setJMSDeliveryTime(  
    long deliveryTime  
) throws JMSEException;  
  
public long getJMSDeliveryTime(  
) throws JMSEException;
```

For more information on the `WLMessage` class, refer to the [weblogic.jms.extensions.WLMessage](#) Javadoc.

## Overriding a Delivery Time

When a producer is created it inherits its `TimeToDeliver` attribute, expressed in milliseconds, from the connection factory used to create the connection that the producer is a part of. Regardless of what time-to-deliver is set on the producer, the destination to which a message is being sent or published can override the setting. An administrator can set the `TimeToDeliverOverride` attribute on a destination in either a relative or scheduled string format.

## Interaction with the Time-to-Live Value

If the specified time-to-live value (`JMSExpiration`) is less than or equal to the specified time-to-deliver value, then the message delivery succeeds. However, the message is then silently expired.

## Setting a Relative Time-to-Deliver Override

A relative `TimeToDeliverOverride` is a String specified as an integer, and is configurable using the Administration Console. For more information, see “[JMS Destinations](#)” in the Administration Console Online Help

## Setting a Scheduled Time-to-Deliver Override

A scheduled `TimeToDeliverOverride` can also be specified using the `weblogic.jms.extensions.schedule` class, which provides methods that take a schedule and return the next scheduled time for delivering messages.

**Table 4-3 Example Time-to-Deliver Schedules**

Example	Description
0 0 0,30 * * * *	Exact next nearest half-hour
* * 0,30 4-5 * * *	Anytime in the first minute of the half hours in the 4 A.M. and 5 A.M. hours
* * * 9-16 * * *	Between 9 A.M. and 5 P.M. (9:00.00 A.M. to 4:59.59 P.M.)
* * * * 8-14 * 2	The second Tuesday of the month
* * * 13-16 * * 0	Between 1 P.M. and 5 P.M. on Sunday
* * * * * 31 *	The last day of the month
* * * * 15 4 1	The next time April 15th occurs on a Sunday
0 0 0 1 * * 2-6;0 0 0 2 * * 1,7	1 A.M. on weekdays; 2 A.M. on weekends

A cron-like string is used to define the schedule. The format is defined by the following BNF syntax:

```
schedule := millisecond second minute hour dayOfMonth month
           dayOfWeek
```

The BNF syntax for specifying the `second` field is as follows:

```
second      := * | secondList
secondList  := secondItem [, secondList]
secondItem  := secondValue | secondRange
SecondRange := secondValue - secondValue
```

Similar BNF statements for milliseconds, minute, hour, day-of-month, month, and day-of-week can be derived from the `second` syntax. The values for each field are defined as non-negative integers in the following ranges:

```
milliSecondValue := 0-999
milliSecondValue := 0-999
secondValue      := 0-59
minuteValue      := 0-59
hourValue        := 0-23
dayOfMonthValue  := 1-31
monthValue       := 1-12
dayOfWeekValue   := 1-7
```

**Note:** These values equate to the same ranges that the `java.util.Calendar` class uses, except for `monthValue`. The `java.util.Calendar` range for `monthValue` is 0-11, rather than 1-12.

Using this syntax, each field can be represented as a range of values indicating all times between the two times. For example, 2-6 in the `dayOfWeek` field indicates Monday through Friday, inclusive. Each field can also be specified as a comma-separated list. For instance, a minute field of 0, 15, 30, 45 means every quarter hour on the quarter hour. Lastly, each field can be defined as both a set of individual values and ranges of values. For example, an hour field of 9-17, 0 indicates between the hours of 9 A.M. and 5 P.M., and on the hour of midnight.

Additional semantics are as follows:

- If multiple schedules are supplied (using a semi-colon (;) as the separator), the next scheduled time for the set is determined using the schedule that returns the soonest value. One use for this is for specifying schedules that change based on the day of the week (see the final example below).
- A value of 1 (one) for the `dayOfWeek` equates to Sunday.
- A value of \* means every time for that field. For instance, a \* in the Month field means every month. A \* in the Hour field means every hour.
- A value of `1` or `last` (not case sensitive) indicates the greatest possible value for a field.
- If a day-of-month is specified that exceeds the normal maximum for a month, then the normal maximum for that month will be specified. For example, if it is February during a leap year and 31 was specified, then the scheduler will schedule as if 29 was specified instead. This means that setting the month field to 31 always indicates the last day of the month.

- If milliseconds are specified, they are rounded down to the nearest 50th of a second. The values are 0, 19, 39, 59, ..., 979, and 999. Thus, 0-40 gets rounded to 0-39 and 50-999 gets rounded to 39-999.

**Note:** When a `Calendar` is not supplied as a method parameter to one of the static methods in this class, the calendar used is a `java.util.GregorianCalendar` with a default `java.util.TimeZone` and a default `java.util.Locale`.

## JMS Schedule Interface

The `weblogic.jms.extensions.schedule` class has methods that will return the next scheduled time that matches the recurring time expression. This expression uses the same syntax as the `TimeToDeliverOverride`. The time returned in milliseconds can be relative or absolute.

For more information on the `WLSession` class, refer to the [weblogic.jms.extensions.Schedule](#) Javadoc.

You can define the next scheduled time after the *given* time using the following method:

```
public static Calendar nextScheduledTime(  
    String schedule,  
    Calendar calendar  
    ) throws ParseException {
```

You can define the next scheduled time after the *current* time using the following method:

```
public static Calendar nextScheduledTime(  
    String schedule,  
    ) throws ParseException {
```

You can define the next scheduled time after the *given* time in absolute milliseconds using the following method:

```
public static long nextScheduledTimeInMillis(  
    String schedule,  
    long timeInMillis  
    ) throws ParseException
```

You can define the next scheduled time after the *given* time in relative milliseconds using the following method:

```
public static long nextScheduledTimeInMillisRelative(  
    String schedule,  
    long timeInMillis  
    ) throws ParseException {
```

You can define the next scheduled time after the *current* time in relative milliseconds using the following method:

```
public static long nextScheduledTimeInMillisRelative(  
    String schedule  
    ) throws ParseException {
```

# Managing Connections

The following sections describe how to manage connections:

- [Defining a Connection Exception Listener](#)
- [Accessing Connection Metadata](#)
- [Starting, Stopping, and Closing a Connection](#)

## Defining a Connection Exception Listener

An exception listener asynchronously notifies an application whenever a problem occurs with a connection. This mechanism is particularly useful for a connection waiting to consume messages that might not be notified otherwise.

**Note:** The purpose of an exception listener is not to monitor all exceptions thrown by a connection, but to deliver those exceptions that would not be otherwise be delivered.

You can define an exception listener for a connection using the following `Connection` method:

```
public void setExceptionListener(  
    ExceptionListener listener  
    ) throws JMSException
```

You must specify an `ExceptionListener` object for the connection.

The JMS Provider notifies an exception listener, if one has been defined, when it encounters a problem with a connection using the following `ExceptionListener` method:

```
public void onException(
    JMSEException exception
)
```

The JMS Provider specifies the exception that describes the problem when calling the method.

You can access the exception listener for a connection using the following `Connection` method:

```
public ExceptionListener getExceptionListener(
) throws JMSEException
```

## Accessing Connection Metadata

You can access the metadata associated with a specific connection using the following `Connection` method:

```
public ConnectionMetaData getMetaData(
) throws JMSEException
```

This method returns a `ConnectionMetaData` object that enables you to access JMS metadata. The following table lists the various type of JMS metadata and the get methods that you can use to access them.

JMS Metadata	Get Method
Version	<code>public String getJMSVersion( ) throws JMSEException</code>
Major version	<code>public int getJMSMajorVersion( ) throws JMSEException</code>
Minor version	<code>public int getJMSMinorVersion( ) throws JMSEException</code>
Provider name	<code>public String getJMSProviderName( ) throws JMSEException</code>

JMS Metadata	Get Method
Provider version	<code>public String getProviderVersion()     ) throws JMSEException</code>
Provider major version	<code>public int getProviderMajorVersion()     ) throws JMSEException</code>
Provider minor version	<code>public int getProviderMinorVersion()     ) throws JMSEException</code>
JMSX property names	<code>public Enumeration getJMSXPropertyNames()     ) throws JMSEException</code>

For more information about the `ConnectionMetaData` class, see the [javax.jms.ConnectionMetaData](#) Javadoc.

## Starting, Stopping, and Closing a Connection

To control the flow of messages, you can start and stop a connection temporarily using the `start()` and `stop()` methods, respectively, as follows.

The `start()` and `stop()` method details are as follows:

```
public void start(  
    ) throws JMSEException  
  
public void stop(  
    ) throws JMSEException
```

A newly created connection is stopped—no messages are received until the connection is started. Typically, other JMS objects are set up to handle messages before the connection is started, as described in “[Setting Up a JMS Application](#)” on page 4-4. Messages may be produced on a stopped connection, but cannot be delivered to a stopped connection.

Once started, you can stop a connection using the `stop()` method. This method performs the following steps:

- Pauses the delivery of all messages. No applications waiting to receive messages will return until the connection is restarted or the time-to-live value associated with the message is reached.

- Waits until all message listeners that are currently processing messages have completed.

Typically, a JMS Provider allocates a significant amount of resources when it creates a connection. When a connection is no longer being used, you should close it to free up resources. A connection can be closed using the following method:

```
public void close(  
    ) throws JMSException
```

This method performs the following steps to execute an orderly shutdown:

- Terminates the receipt of all pending messages. Applications may return a message or null if a message was not available at the time of the close.
- Waits until all message listeners that are currently processing messages have completed.
- Rolls back in-process transactions on its transacted sessions (unless such transactions are part of an external JTA user transaction). For more information about JTA user transactions, see [“Using JTA User Transactions” on page 5-5](#).
- Does not force an acknowledge of client-acknowledged sessions. By not forcing an acknowledge, no messages are lost for queues and durable subscriptions that require reliable processing.

When you close a connection, all associated objects are also closed. You can continue to use the message objects created or received via the connection, except the received message’s `acknowledge()` method. Closing a closed connection has no effect.

**Note:** Attempting to acknowledge a received message from a closed connection’s session throws an `IllegalStateException`.

## Managing Sessions

The following sections describe how to manage sessions, including:

- [Defining a Session Exception Listener](#)
- [Closing a Session](#)

# Defining a Session Exception Listener

An exception listener asynchronously notifies a client in the event a problem occurs with a session. This is particularly useful for a session waiting to consume messages that might not be notified otherwise.

**Note:** The purpose of an exception listener is not to monitor all exceptions thrown by a session, only to deliver those exceptions that would otherwise be undelivered.

You can define an exception listener for a session using the following `WLSession` method:

```
public void setExceptionListener(  
    ExceptionListener listener  
) throws JMSEException
```

You must specify an `ExceptionListener` object for the session.

The JMS Provider notifies an exception listener, if one has been defined, when it encounters a problem with a session using the following `ExceptionListener` method:

```
public void onException(  
    JMSEException exception  
)
```

The JMS Provider specifies the exception encountered that describes the problem when calling the method.

You can access the exception listener for a session using the following `WLSession` method:

```
public ExceptionListener getExceptionListener(  
) throws JMSEException
```

**Note:** Because there can only be one thread per session, an exception listener and message listener (used for asynchronous message delivery) cannot execute simultaneously. Consequently, if a message listener is executing at the time a problem occurs, execution of the exception listener is blocked until the message listener completes its execution. For more information about message listeners, see [“Receiving Messages Asynchronously” on page 4-30](#).

## Closing a Session

As with connections, a JMS Provider allocates a significant amount of resources when it creates a session. When a session is no longer being used, it is recommended that it be closed to free up resources. A session can be closed using the following `Session` method:

```
public void close(  
    ) throws JMSEException
```

**Note:** The `close()` method is the only `Session` method that can be invoked from a thread that is separate from the session thread.

This method performs the following steps to execute an orderly shutdown:

- Terminates the receipt of all pending messages. Applications may return a message or null if a message was not available at the time of the close.
- Waits until all message listeners that are currently processing messages have completed.
- Rolls back in-process transactions (unless such transactions are part of external JTA user transaction). For more information about JTA user transactions, see [“Using JTA User Transactions” on page 5-5](#).
- Does not force an acknowledge of client acknowledged sessions, ensuring that no messages are lost for queues and durable subscriptions that require reliable processing.

When you close a session, all associated producers and consumers are also closed.

**Note:** If you wish to issue the `close()` method within an `onMessage()` method call, the system administrator must select the `Allow Close In OnMessage` check box when configuring the connection factory. For more information, see [“JMS Connection Factories”](#) in the Administration Console Online Help.

## Creating Destinations Dynamically

You can create destinations dynamically using:

- `weblogic.jms.extensions.JMSHelper` class methods
- Temporary destinations

The associated procedures for creating dynamic destinations are described in the following sections.

### Using the `JMSHelper` Class Methods

You can dynamically submit an asynchronous request to create a queue or topic, respectively, using the following `JMSHelper` methods:

```
static public void createPermanentQueueAsync(  
    Context ctx,  
    String jmsServerName,  
    String queueName,  
    String jndiName  
) throws JMSEException  
  
static public void createPermanentTopicAsync(  
    Context ctx,  
    String jmsServerName,  
    String topicName,  
    String jndiName  
) throws JMSEException
```

You must specify the JNDI initial context, name of the JMS server to be associated with the destination, name of the destination (queue or topic), and name used to look up the destination within the JNDI namespace.

Each method updates the following:

- Configuration file associated with the specified domain to include the dynamically created destination
- JNDI namespace to advertise the destination

**Note:** Either method call can fail without throwing an exception. In addition, a thrown exception does not necessarily indicate that the method call failed.

The time required to create the destination on the JMS server and propagate the information to the JNDI namespace can be significant. The propagation delay increases if the environment contains multiple servers. It is recommended that you test

for the existence of the queue or topic, respectively, using the session `createQueue()` or `createTopic()` method, rather than perform a JNDI lookup. By doing so, you can avoid some of the propagation-specific delay.

For example, the following method, `findQueue()`, attempts to access a dynamically created queue, and if unsuccessful, sleeps for a specified interval before retrying. A maximum retry count is established to prevent an infinite loop.

```
private static Queue findQueue (
    QueueSession queueSession,
    String jmsServerName,
    String queueName,
    int retryCount,
    long retryInterval
) throws JMSEException
{
    String wlsQueueName = jmsServerName + "/" + queueName;
    String command = "QueueSession.createQueue(" +
        wlsQueueName + ")";
    long startTimeMillis = System.currentTimeMillis();
    for (int i=retryCount; i>=0; i--) {
        try {
            System.out.println("Trying " + command);
            Queue queue = queueSession.createQueue(wlsQueueName);
            System.out.println(command + "succeeded after " +
                (retryCount - i + 1) + " tries in " +
                (System.currentTimeMillis() - startTimeMillis) +
                " millis.");
            return queue;
        } catch (JMSEException je) {
            if (retryCount == 0) throw je;
        }
        try {
            System.out.println(command + "> failed, pausing " +
                retryInterval + " millis.");
            Thread.sleep(retryInterval);
        } catch (InterruptedException ignore) {}
    }
    throw new JMSEException("out of retries");
}
```

You can then call the `findQueue()` method after the `JMSHelper` class method call to retrieve the dynamically created queue once it becomes available. For example:

```
JMSHelper.createPermanentQueueAsync(ctx, domain, jmsServerName,
    queueName, jndiName);
```

```
Queue queue = findQueue(qsess, jmsServerName, queueName,  
    retry_count, retry_interval);
```

For more information on the `JMSHelper` class, refer to the [weblogic.jms.extensions.JMSHelper](#) Javadoc.

## Using Temporary Destinations

Temporary destinations enable an application to create a destination, as required, without the system administration overhead associated with configuring and creating a server-defined destination.

The WebLogic JMS server can use the `JMSReplyTo` header field to return a response to the application. The application may optionally set the `JMSReplyTo` header field of its messages to its temporary destination name to advertise the temporary destination that it is using to other applications.

Temporary destinations exist only for the duration of the current connection, unless they are removed using the `delete()` method, described in “[Deleting a Temporary Destination](#)” on page 4-53.

Because messages are never available if the server is restarted, all `PERSISTENT` messages are silently made `NON_PERSISTENT`. As a result, temporary destinations are not suitable for business logic that must survive a restart.

**Note:** Before creating a temporary destination (queue or topic), you must use the Administration Console to configure the JMS server to use temporary destinations. This is done by using the JMS server’s `Temporary Template` attribute to select a JMS template that is configured in the same domain. For more information about configuring a JMS server, see “[JMS Server](#)” in the Administration Console Online Help.

The following sections describe how to create a temporary queue (PTP) or temporary topic (Pub/Sub).

### Creating a Temporary Queue

You can create a temporary queue using the following `QueueSession` method:

```
public TemporaryQueue createTemporaryQueue(  
    ) throws JMSException
```

For example, to create a reference to a `TemporaryQueue` that will exist only for the duration of the current connection, use the following method call:

```
QueueSender = Session.createTemporaryQueue();
```

### Creating a Temporary Topic

You can create a temporary topic using the following `TopicSession` method:

```
public TemporaryTopic createTemporaryTopic(  
    ) throws JMSEException
```

For example, to create a reference to a temporary topic that will exist only for the duration of the current connection, use the following method call:

```
TopicPublisher = Session.createTemporaryTopic();
```

### Deleting a Temporary Destination

When you finish using a temporary destination, you can delete it (to release associated resources) using the following `TemporaryQueue` or `TemporaryTopic` method:

```
public void delete(  
    ) throws JMSEException
```

## Setting Up Durable Subscriptions

WebLogic JMS supports durable and non-durable subscriptions.

For durable subscriptions, WebLogic JMS stores a message in a persistent file or database until the message has been delivered to the subscribers or has expired, even if those subscribers are not *active* at the time that the message is delivered. A subscriber is considered active if the Java object that represents it exists. Durable subscriptions are supported for Pub/Sub messaging only.

**Note:** Durable subscriptions cannot be created for distributed topics. However, you can still create a durable subscription on distributed topic member and the other topic members will forward the messages to the member that has the durable subscription. For more information on using distributed topics, see [“Using Distributed Destinations”](#) on page 4-88.

For non-durable subscriptions, WebLogic JMS delivers messages only to applications with an active session. Messages sent to a topic while an application is not listening are never delivered to that application. In other words, non-durable subscriptions last only as long as their subscriber objects. By default, subscribers are non-durable.

The following sections describe:

- [Defining the Persistent Store](#)
- [Defining the Client ID](#)
- [Creating Subscribers for a Durable Subscription](#)
- [Deleting Durable Subscriptions](#)
- [Modifying Durable Subscriptions](#)
- [Managing Durable Subscriptions](#)

## Defining the Persistent Store

You must configure a persistent file or database store and assign it to your JMS server so WebLogic JMS can store a message until it has been delivered to the subscribers or has expired.

- Create a JMS file store or JMS JDBC backing store using the Stores node.
- Target the configured store to your JMS server by selecting it from the Store field’s drop-down list on the JMS Server → Configuration → General tab.

**Note:** No two JMS servers can use the same backing store.

For more information about configuring JMS stores, see [“Managing JMS”](#) in the *Administration Guide*.

## Defining the Client ID

To support durable subscriptions, a client identifier (client ID) must be defined for the connection.

**Note:** The JMS client ID is not necessarily equivalent to the WebLogic Server username, that is, a name used to authenticate a user in the WebLogic security realm. You can, of course, set the JMS client ID to the WebLogic Server username, if it is appropriate for your JMS application.

The client ID can be supplied in two ways:

- The first method is to configure the connection factory with the client ID. For WebLogic JMS, this means adding a separate connection factory definition during configuration for each client ID. Applications then look up their own topic connection factories in JNDI and use them to create connections containing their own client IDs. For more information about configuring a connection factory with a client ID, see “[JMS Connection Factories](#)” in the Administration Console Online Help.
- Alternatively, the preferred method is for an application that can set its client ID in the connection after the connection is created by calling the following connection method:

```
public void setClientID(  
    String clientID  
) throws JMSEException
```

You must specify a unique client ID. If you use this alternative approach, you can use the default connection factory (if it is acceptable for your application) and avoid the need to modify the configuration information. However, applications with durable subscriptions must ensure that they call `setClientID()` *immediately after* creating their topic connection. For information on the default connection factory, see “[Managing JMS](#)” in the *Administration Guide*.

If a client ID is already defined for the connection, an `IllegalStateException` is thrown. If the specified client ID is already defined for another connection, an `InvalidClientIDException` is thrown.

**Note:** When specifying the client ID using the `setClientID()` method, there is a risk that a duplicate client ID may be specified without throwing an exception. For example, if the client IDs for two separate connections are

set simultaneously to the same value, a race condition may occur and the same value may be assigned to both connections. You can avoid this risk of duplication by specifying the client ID during configuration.

To display a client ID and test whether or not a client ID has already been defined, use the following Connection method:

```
public String getClientID(  
    ) throws JMSEException
```

**Note:** Support for durable subscriptions is a feature unique to the Pub/Sub messaging model, so client IDs are used only with topic connections; queue connections also contain client IDs, but JMS does not use them.

Durable subscriptions should not be created for a temporary topic, because a temporary topic is designed to exist only for the duration of the current connection.

## Creating Subscribers for a Durable Subscription

You can create subscribers for a durable subscription using the following `TopicSession` methods:

```
public TopicSubscriber createDurableSubscriber(  
    Topic topic,  
    String name  
    ) throws JMSEException  
  
public TopicSubscriber createDurableSubscriber(  
    Topic topic,  
    String name,  
    String messageSelector,  
    boolean noLocal  
    ) throws JMSEException
```

You must specify the name of the topic for which you are creating a subscriber, and the name of the durable subscription. You may also specify a message selector for filtering messages and a `noLocal` flag (described later in this section). Message selectors are described in more detail in [“Filtering Messages” on page 4-66](#). If you do not specify a `messageSelector`, by default all messages are searched.

An application can use a JMS connection to both publish and subscribe to the same topic. Because topic messages are delivered to all subscribers, an application can receive messages it has published itself. To prevent this, a JMS application can set a `noLocal` flag to `true`. The `noLocal` value defaults to `false`.

The durable subscription `name` must be unique per client ID. For information on defining the client ID for the connection, see [“Defining the Client ID” on page 4-55](#).

Only one session can define a subscriber for a particular durable subscription at any given time. Multiple subscribers can access the durable subscription, but not at the same time. Durable subscriptions are stored within the file or database.

## Deleting Durable Subscriptions

To delete a durable subscription, you use the following `TopicSession` method:

```
public void unsubscribe(  
    String name  
) throws JMSException
```

You must specify the name of the durable subscription to be deleted.

You cannot delete a durable subscription if any of the following are true:

- A `TopicSubscriber` is still active on the session.
- A message received by the durable subscription is part of a transaction or has not yet been acknowledged in the session.

**Note:** You can also delete durable subscriptions from the Administration Console. For information on managing durable subscriptions, see [“Managing Durable Subscriptions” on page 4-58](#).

## Modifying Durable Subscriptions

To modify a durable subscription, perform the following steps:

1. Optionally, delete the durable subscription, as described in [“Deleting Durable Subscriptions” on page 4-57](#).

This step is optional. If not explicitly performed, the deletion will be executed implicitly when the durable subscription is recreated in the next step.

2. Use the methods described in “[Creating Subscribers for a Durable Subscription](#)” on page 4-56 to recreate a durable subscription of the same name, but specifying a different topic name, message selector, or `noLocal` value.

The durable subscription is recreated based on the new values.

**Note:** When recreating a durable subscription, be careful to avoid creating a durable subscription with a duplicate name. For example, if you attempt to delete a durable subscription from a JMS server that is unavailable, the delete call fails. If you subsequently create a durable subscription with the same name on a different JMS server, you may experience unexpected results when the first JMS server becomes available. Because the original durable subscription has not been deleted, when the first JMS server again becomes available, there will be two durable subscriptions with duplicate names.

## Managing Durable Subscriptions

You can monitor and delete durable subscriptions from the Administration Console. For more information, see “[Managing JMS](#)” in the *Administration Guide*.

## Setting and Browsing Message Header and Property Fields

WebLogic JMS provides a set of standard header fields that you can define to identify and route messages. In addition, property fields enable you to include application-specific header fields within a message, extending the standard set. You can use the message header and property fields to convey information between communicating processes.

The primary reason for including data in a property field rather than in the message body is to support message filtering via message selectors. Except for XML message extensions, data in the message body cannot be accessed via message selectors. For

example, suppose you use a property field to assign high priority to a message. You can then design a message consumer containing a message selector that accesses this property field and selects only messages of expedited priority. For more information about selectors, see [“Filtering Messages” on page 4-66](#).

## Setting Message Header Fields

JMS messages contain a standard set of header fields that are always transmitted with the message. They are available to message consumers that receive messages, and some fields can be set by the message producers that send messages. Once a message is received, its header field values can be modified.

For a description of the standard messages header fields, see [“Message Header Fields” on page 2-16](#).

The following table lists the Message class set and get methods for each of the supported data types.

**Note:** In some cases, the `send()` method overrides the header field value set using the `set()` method, as indicated in the following table.

Header Field	Set Method	Get Method
JMSCorrelationID	<pre>public void setJMSCorrelationID(     String correlationID ) throws JMSEException</pre>	<pre>public String getJMSCorrelationID( ) throws JMSEException  public byte[] getJMSCorrelationIDAsBytes( ) throws JMSEException</pre>
JMSDestination <sup>1</sup>	<pre>public void setJMSDestination(     Destination destination ) throws JMSEException</pre>	<pre>public Destination getJMSDestination( ) throws JMSEException</pre>
JMSDeliveryMode <sup>1</sup>	<pre>public void setJMSDeliveryMode(     int deliveryMode ) throws JMSEException</pre>	<pre>public int getJMSDeliveryMode( ) throws JMSEException</pre>

## 4 Developing a WebLogic JMS Application

---

Header Field	Set Method	Get Method
JMSDeliveryTime <sup>1</sup>	<pre>public void setJMSDeliveryTime(     long deliveryTime ) throws JMSEException</pre>	<pre>public long getJMSDeliveryTime( ) throws JMSEException</pre>
JMSDeliveryMode <sup>1</sup>	<pre>public void setJMSDeliveryMode(     int deliveryMode ) throws JMSEException</pre>	<pre>public int getJMSDeliveryMode( ) throws JMSEException</pre>
JMSMessageID <sup>1</sup>	<pre>public void setJMSMessageID(     String id ) throws JMSEException</pre>	<pre>public String getJMSMessageID( ) throws JMSEException</pre>
	<p>In addition to the set method, the <code>weblogic.jms.extensions.JMSHelper</code> class provides the following methods to convert between pre-WebLogic JMS 6.0 and 6.1 JMSMessageID formats:</p> <pre>public void oldJMSMessageIDToNew(     String id,     long timeStamp ) throws JMSEException  public void newJMSMessageIDToOld(     String id,     long timeStamp ) throws JMSEException</pre>	
JMSPriority <sup>1</sup>	<pre>public void setJMSPriority(     int priority ) throws JMSEException</pre>	<pre>public int getJMSPriority( ) throws JMSEException</pre>
JMSRedelivered <sup>1</sup>	<pre>public void setJMSRedelivered(     boolean redelivered ) throws JMSEException</pre>	<pre>public boolean getJMSRedelivered( ) throws JMSEException</pre>
JMSRedeliveryLimit <sup>1</sup>	<pre>public void setJMSRedeliveryLimit(     int redelivered ) throws JMSEException</pre>	<pre>public int getJMSRedeliveryLimit( ) throws JMSEException</pre>

## Setting and Browsing Message Header and Property Fields

---

Header Field	Set Method	Get Method
JMSReplyTo	<code>public void setJMSReplyTo(Destination replyTo) throws JMSEException</code>	<code>public Destination getJMSReplyTo() throws JMSEException</code>
JMSTimeStamp <sup>1</sup>	<code>public void setJMSTimeStamp(long timestamp) throws JMSEException</code>	<code>public long getJMSTimeStamp() throws JMSEException</code>
JMSType	<code>public void setJMSType(String type) throws JMSEException</code>	<code>public String getJMSType() throws JMSEException</code>

1. The corresponding `set()` method has no impact on the message header field when the `send()` method is executed. If set, this header field value will be overridden during the `send()` operation.

The `examples.jms.sender.SenderServlet` example, provided with WebLogic Server in the `WL_HOME\samples\server\src\examples\jms\sender` directory, where `WL_HOME` is the top-level directory of your WebLogic Platform installation, shows how to set header fields in messages that you send and how to display message header fields after they are sent.

For example, the following code, which appears after the `send()` method, displays the message ID that was assigned to the message by WebLogic JMS:

```
System.out.println("Sent message " +
    msg.getJMSMessageID() + " to " +
    msg.getJMSDestination());
```

## Setting Message Property Fields

To set a property field, call the appropriate set method and specify the property name and value. To read a property field, call the appropriate get method and specify the property name.

The sending application can set properties in the message, and the receiving application can subsequently view them. The receiving application cannot change the properties without first clearing them using the following `clearProperties()` method:

## 4 Developing a WebLogic JMS Application

```
public void clearProperties(  
    ) throws JMSEException
```

This method does not clear the message header fields or body.

**Note:** The `JMSX` property name prefix is reserved for JMS. The connection metadata contains a list of `JMSX` properties, which can be accessed as an enumerated list using the `getJMSXPropertyNames()` method. For more information, see [“Accessing Connection Metadata” on page 4-45](#).

The `JMS_` property name prefix is reserved for provider-specific properties; it is not intended for use with standard JMS messaging.

The property field can be set to any of the following types: boolean, byte, double, float, int, long, short, or string. The following table lists the Message class set and get methods for each of the supported data types.

**Table 4-4 Message Property Set and Get Methods for Data Types**

Data Type	Set Method	Get Method
boolean	<pre>public void setBooleanProperty(     String name,     boolean value     ) throws JMSEException</pre>	<pre>public boolean getBooleanProperty(     String name     ) throws JMSEException</pre>
byte	<pre>public void setByteProperty(     String name,     byte value     ) throws JMSEException</pre>	<pre>public byte getByteProperty(     String name     ) throws JMSEException</pre>
double	<pre>public void setDoubleProperty(     String name,     double value     ) throws JMSEException</pre>	<pre>public double getDoubleProperty(     String name     ) throws JMSEException</pre>
float	<pre>public void setFloatProperty(     String name,     float value     ) throws JMSEException</pre>	<pre>public float getFloatProperty(     String name     ) throws JMSEException</pre>
int	<pre>public void setIntProperty(     String name,     int value     ) throws JMSEException</pre>	<pre>public int getIntProperty(     String name     ) throws JMSEException</pre>

**Table 4-4 Message Property Set and Get Methods for Data Types (Continued)**

Data Type	Set Method	Get Method
long	public void setLongProperty( String name, long value) throws JMSEException	public long getLongProperty( String name ) throws JMSEException
short	public void setShortProperty( String name, short value ) throws JMSEException	public short getShortProperty( String name ) throws JMSEException
String	public void setStringProperty( String name, String value ) throws JMSEException	public String getStringProperty( String name ) throws JMSEException

In addition to the set and get methods described in the previous table, you can use the `setObjectProperty()` and `getObjectProperty()` methods to use the objectified primitive values of the property type. When the objectified value is used, the property type can be determined at execution time rather than during the compilation. The valid object types are boolean, byte, double, float, int, long, short, and string.

You can access all property field names using the following Message method:

```
public Enumeration getPropertyNames(  
    ) throws JMSEException
```

This method returns all property field names as an enumeration. You can then retrieve the value of each property field by passing the property field name to the appropriate get method, as described in the previous table, based on the property field data type.

The following table is a conversion chart for message properties. It allows you to identify the type that can be read based on the type that has been written.

**Table 4-5 Message Property Conversion Chart**

Property Written As. . .	Can Be Read As. . .							
	boolea n	byte	doubl e	float	int	long	short	Strin g
boolean	X							X

**Table 4-5 Message Property Conversion Chart (Continued)**

Property Written As. . .	Can Be Read As. . .							
	boolean	byte	double	float	int	long	short	String
byte		X			X	X	X	X
double			X					X
float			X	X				X
int					X	X		X
long						X		X
Object	X	X	X	X	X	X	X	X
short					X	X	X	X
String	X	X	X	X	X	X	X	X

You can test whether or not a property value has been set using the following `Message` method:

```
public boolean propertyExists(
    String name
) throws JMSEException
```

You specify a property name and the method returns a boolean value indicating whether or not the property exists.

For example, the following code sets two `String` properties and an `int` property:

```
msg.setStringProperty("User", user);
msg.setStringProperty("Category", category);
msg.setIntProperty("Rating", rating);
```

For more information about message property fields, see [“Message Property Fields” on page 2-20](#) or the `javax.jms.Message` Javadoc.

## Browsing Header and Property Fields

**Note:** Only queue message header and property fields can be browsed. You cannot browse topic message header and property fields.

You can browse the header and property fields of messages on a queue using the following `QueueSession` methods:

```
public QueueBrowser createBrowser(  
    Queue queue  
    ) throws JMSException  
  
public QueueBrowser createBrowser(  
    Queue queue,  
    String messageSelector  
    ) throws JMSException
```

You must specify the queue that you wish to browse. You may also specify a message selector to filter messages that you are browsing. Message selectors are described in more detail in [“Filtering Messages” on page 4-66](#).

Once you have defined a queue, you can access the queue name and message selector associated with a queue browser using the following `QueueBrowser` methods:

```
public Queue getQueue(  
    ) throws JMSException  
  
public String getMessageSelector(  
    ) throws JMSException
```

In addition, you can access an enumeration for browsing the messages using the following `QueueBrowser` method:

```
public Enumeration getEnumeration(  
    ) throws JMSException
```

The `examples.jms.queue.QueueBrowser` example, provided with WebLogic Server in the `WL_HOME\samples\server\src\examples\jms\queue` directory, where `WL_HOME` is the top-level directory of your WebLogic Platform installation, shows how to access the header fields of received messages.

For example, the following code line is an excerpt from the `QueueBrowser` example and creates the `QueueBrowser` object:

```
qbrowser = qsession.createBrowser(queue);
```

The following provides an excerpt from the `displayQueue()` method defined in the `QueueBrowser` example. In this example, the `QueueBrowser` object is used to obtain an enumeration that is subsequently used to scan the queue's messages.

```
public void displayQueue(
    ) throws JMSException
{
    Enumeration e = qbrowser.getEnumeration();
    Message m = null;

    if (! e.hasMoreElements()) {
        System.out.println("There are no messages on this queue.");
    } else {

        System.out.println("Queued JMS Messages: ");
        while (e.hasMoreElements()) {
            m = (Message) e.nextElement();
            System.out.println("Message ID " + m.getJMSMessageID() +
                " delivered " + new Date(m.getJMSTimestamp())
                " to " + m.getJMSDestination());
        }
    }
}
```

When a queue browser is no longer being used, you should close it to free up resources. For more information, see [“Releasing Object Resources” on page 4-34](#).

For more information about the `QueueBrowser` class, see the [javax.jms.QueueBrowser Javadoc](#).

## Filtering Messages

In many cases, an application does not need to be notified of every message that is delivered to it. Message selectors can be used to filter unwanted messages, and subsequently improve performance by minimizing their impact on network traffic.

Message selectors operate as follows:

- The sending application sets message header or property fields to describe or classify a message in a standardized way.
- The receiving applications specify a simple query string to filter the messages that they want to receive.

Because message selectors cannot reference the contents (body) of a message, some information may be duplicated in the message property fields (except in the case of XML messages).

You specify a selector when creating a queue receiver or topic subscriber, as an argument to the `QueueSession.createReceiver()` or `TopicSession.createSubscriber()` methods, respectively. For information about creating queue receivers and topic subscribers, see “[Step 5: Create Message Producers and Message Consumers Using the Session and Destinations](#)” on page 4-11.

The following sections describe how to define a message selector using SQL statements and XML selector methods, and how to update message selectors. For more information about setting header and property fields, see “[Setting and Browsing Message Header and Property Fields](#)” on page 4-58 and “[Setting Message Property Fields](#)” on page 4-61, respectively.

## Defining Message Selectors Using SQL Statements

A message selector is a boolean expression. It consists of a String with a syntax similar to the `where` clause of an SQL `select` statement.

The following excerpts provide examples of selector expressions.

```
salary > 64000 and dept in ('eng','qa')

(product like 'WebLogic%' or product like '%T3')
and version > 3.0

hireyear between 1990 and 1992
or fireyear is not null

fireyear - hireyear > 4
```

The following example shows how to set a selector when creating a queue receiver that filters out messages with a priority lower than 6.

```
String selector = "JMSPriority >= 6";
qsession.createReceiver(queue, selector);
```

The following example shows how to set the same selector when creating a topic subscriber.

```
String selector = "JMSPriority >= 6";
qsession.createSubscriber(topic, selector);
```

For more information about the message selector syntax, see the [javax.jms.Message](#) Javadoc.

# Defining XML Message Selectors Using XML Selector Method

For XML message types, in addition to using the SQL selector expressions described in the previous section to define message selectors, you can use the following method:

```
String JMS_BEA_SELECT(String type, String expression)
```

`JMS_BEA_SELECT` is a built-in function in WebLogic JMS SQL syntax. You specify the syntax type, which must be set to `xpath` (XML Path Language) and an XPath expression. The XML path language is defined in the XML Path Language (XPath) document, which is available at the XML Path Language Web site at:

<http://www.w3.org/TR/xpath>

**Note:** Pay careful attention to your XML message syntax, since malformed XML messages (for example, a missing end tag) will not match any XML selector.

The method returns a null value under the following circumstances:

- The message does not parse.
- The message parses, but the element is not present.
- If a message parses and the element is present, but the message contains no value (for example, `<order></order>`).

For example, consider the following XML excerpt:

```
<order>
  <item>
    <id>007</id>
    <name>Hand-held Power Drill</name>
    <description>Compact, assorted colors.</description>
    <price>$34.99</price>
  </item>
  <item>
    <id>123</id>
    <name>Mitre Saw</name>
    <description>Three blades sizes.</description>
    <price>$69.99</price>
  </item>
</order>
```

```

    </item>
    <item>
        <id>66</id>
        <name>Socket Wrench Set</name>
        <description>Set of 10.</description>
        <price>$19.99</price>
    </item>
</order>

```

The following example shows how to retrieve the name of the second item in the previous example. This method call returns the string, `Mitre Saw`.

```
String sel = "JMS_BEA_SELECT('xpath',
    '/order/item[2]/name/text()') = 'Mitre Saw'";
```

Pay careful attention to the use of double and single quotes and spaces. Note the use of single quotes around `xpath`, the XML tag, and the string value.

The following example shows how to retrieve the ID of the third item in the previous example. This method call returns the string, `66`.

```
String sel = "JMS_BEA_SELECT('xpath',
    '/order/item[3]/id/text()') = '66'";
```

## Displaying Message Selectors

You can use the following `MessageConsumer` method to display a message selector:

```
public String getMessageSelector(
) throws JMSEException
```

This method returns either the currently defined message selector or null if a message selector is not defined.

## Indexing Topic Subscriber Message Selectors To Optimize Performance

For a certain class of applications, WebLogic JMS can significantly optimize topic subscriber message selectors by indexing them. These applications typically have a large number of subscribers, each with a unique identifier (like a user name), and they need to be able to quickly send a message to a single subscriber, or to a list of

subscribers. A typical example is an instant messaging application where each subscriber corresponds to a different user, and each message contains a list of one or more target users.

To activate optimized subscriber message selectors, subscribers must use the following syntax for their selectors:

```
"identifier IS NOT NULL"
```

where *identifier* is an arbitrary string that is not a predefined JMS message property (e.g., neither `JMSCorrelationID` nor `JMSType`). Multiple subscribers can share the same identifier.

WebLogic JMS uses this exact message selector syntax as a hint to build internal subscriber indexes. Message selectors that do not follow the syntax, or that include additional `OR` and `AND` clauses, are still honored, but do not activate the optimization.

Once subscribers have registered using this message selector syntax, a message published to the topic can target specific subscribers by including one or more identifiers in the message's user properties, as illustrated in the following example:

```
// Set up a named subscriber, where "wilma" is the name of
// the subscriber and subscriberSession is a JMS TopicSession.
// Note that the selector syntax used activates the optimization.

TopicSubscriber topicSubscriber =
    subscriberSession.createSubscriber(
        (Topic)context.lookup("IMTopic"),
        "Wilma IS NOT NULL",
        /* noLocal= */ true);

// Send a message to subscribers "Fred" and "Wilma",
// where publisherSession is a JMS TopicSession. Subscribers
// with message selector expressions "Wilma IS NOT NULL"
// or "Fred IS NOT NULL" will receive this message.

TopicPublisher topicPublisher =
    publisherSession.createPublisher(
        (Topic)context.lookup("IMTopic");

TextMessage msg =
    publisherSession.createTextMessage("Hi there!");
msg.setBooleanProperty("Fred", true);
msg.setBooleanProperty("Wilma", true);

topicPublisher.publish(msg);
```

**Notes:**

The optimized message selector and message syntax is based on the standard JMS API; therefore, applications that use this syntax will also work on versions of WebLogic JMS that do not have optimized message selectors, as well as on non-WebLogic JMS products. However, these versions will not perform as well as versions that include this enhancement.

The message selector optimization will have no effect on applications that use the `MULTICAST_NO_ACKNOWLEDGE` acknowledge mode. These applications have no need for the enhancement anyway, since the message selection occurs on the client side rather than the server side.

## Defining Server Session Pools

**Note:** Session pools are now used rarely, as they are not a required part of the J2EE specification, do not support JTA user transactions, and are largely superseded by message-driven beans (MDBs), which are simpler, easier to manage, and more capable. For more information on designing MDBs, see “[Designing and Developing Message-Driven Beans](#)” in *Programming WebLogic Enterprise JavaBeans*.

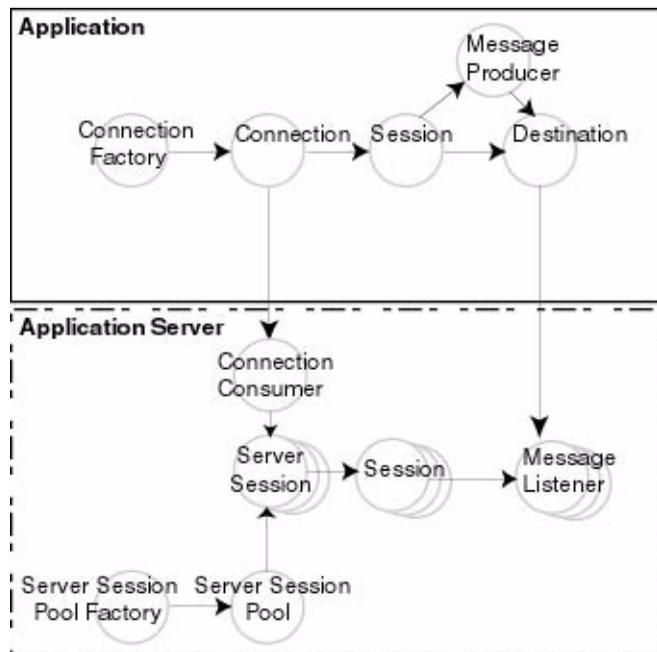
WebLogic JMS implements an optional JMS facility for defining a server-managed pool of server sessions. This facility enables an application to process messages concurrently.

The server session pool:

- Receives messages from a destination and passes them to a server-side message listener that you provide to process messages. The message listener class provides an `onMessage()` method that processes a message.
- Processes messages in parallel by managing a pool of JMS sessions, each of which executes a single-threaded `onMessage()` method.

The following figure illustrates the server session pool facility, and the relationship between the application and the application server components.

Figure 4-3 Server Session Pool Facility

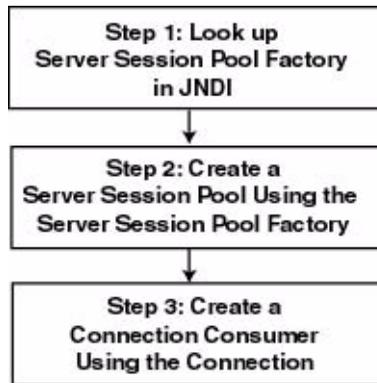


As illustrated in the figure, the application provides a single-threaded message listener. The connection consumer, implemented by JMS on the application server, performs the following tasks to process one or more messages:

1. Gets a server session from the server session pool.
2. Gets the server session's session.
3. Loads the session with one or more messages.
4. Starts the server session to consume messages.
5. Releases the server session back to pool when finished processing messages.

The following figure illustrates the steps required to prepare for concurrent message processing.

Figure 4-4 Preparing for Concurrent Message Processing



Applications can use other application server providers' session pool implementations within this flow. Server session pools can also be implemented using Message Driven Beans. For information on using message driven beans to implement server session pools, see [Programming WebLogic Enterprise JavaBeans](#).

If the session pool and connection consumer were defined during configuration, you can skip this section. For more information on configuring server session pools and connection consumers, see [Managing JMS](#) in the *Administration Guide*.

Currently, WebLogic JMS does *not* support the optional `TopicConnection.createDurableConnectionConsumer()` operation. For more information on this advanced JMS operation, refer to [Sun Microsystems' JMS Specification](#).

## Step 1: Look Up Server Session Pool Factory in JNDI

You use a server session pool factory to create a server session pool.

WebLogic JMS defines one `ServerSessionPoolFactory` object, by default: `weblogic.jms.ServerSessionPoolFactory:<name>`, where `<name>` specifies the name of the JMS server to which the session pool is created.

Once it has been configured, you can look up a server session pool factory by first establishing a JNDI context (`context`) using the `NamingManager.InitialContext()` method. For any application other than a servlet application, you must pass an environment used to create the initial context. For more information, see the `NamingManager.InitialContext()` Javadoc.

Once the context is defined, to look up a server session pool factory in JNDI use the following code:

```
factory = (ServerSessionPoolFactory) context.lookup(<ssp_name>);
```

The `<ssp_name>` specifies a qualified or non-qualified server session pool factory name.

For more information about server session pool factories, see “[ServerSessionPoolFactory](#)” on page 2-22 or the `weblogic.jms.ServerSessionPoolFactory` Javadoc.

## Step 2: Create a Server Session Pool Using the Server Session Pool Factory

You can create a server session pool for use by queue (PTP) or topic (Pub/Sub) connection consumers, using the `ServerSessionPoolFactory` methods described in the following sections.

For more information about server session pools, see “[ServerSessionPool](#)” on page 2-22 or the `javax.jms.ServerSessionPool` Javadoc.

### Create a Server Session Pool for Queue Connection Consumers

The `ServerSessionPoolFactory` provides the following method for creating a server session pool for queue connection consumers:

```
public ServerSessionPool getServerSessionPool(  
    QueueConnection connection,  
    int maxSessions,  
    boolean transacted,  
    int ackMode,  
    String listenerClassName  
) throws JMSEException
```

You must specify the queue connection associated with the server session pool, the maximum number of concurrent sessions that can be retrieved by the connection consumer (to be created in step 3), whether or not the sessions are transacted, the acknowledge mode (applicable for non-transacted sessions only), and the message listener class that is instantiated and used to receive and process messages concurrently.

For more information about the `ServerSessionPoolFactory` class methods, see the [weblogic.jms.ServerSessionPoolFactory](#) Javadoc. For more information about the `ConnectionConsumer` class, see the [javax.jms.ConnectionConsumer](#) Javadoc.

### Create a Server Session Pool for Topic Connection Consumers

The `ServerSessionPoolFactory` provides the following method for creating a server session pool for topic connection consumers:

```
public ServerSessionPool getServerSessionPool(
    TopicConnection connection,
    int maxSessions,
    boolean transacted,
    int ackMode,
    String listenerClassName
) throws JMSEException
```

You must specify the topic connection associated with the server session pool, the maximum number of concurrent sessions that can be retrieved by the connection (to be created in step 3), whether or not the sessions are transacted, the acknowledge mode (applicable for non-transacted sessions only), and the message listener class that is instantiated and used to receive and process messages concurrently.

For more information about the `ServerSessionPoolFactory` class methods, see the [weblogic.jms.ServerSessionPoolFactory](#) Javadoc. For more information about the `ConnectionConsumer` class, see the [javax.jms.ConnectionConsumer](#) Javadoc.

## Step 3: Create a Connection Consumer

You can create a connection consumer for retrieving server sessions and processing messages concurrently using one of the following methods:

- Configuring the server session pool and connection consumer during the configuration, as described in the “[Managing JMS](#)” in the *Administration Guide*
- Including in your application the Connection methods described in the following sections

For more information about the `ConnectionConsumer` class, see “[ConnectionConsumer](#)” on page 2-23 or the `javax.jms.ConnectionConsumer` Javadoc.

### Create a Connection Consumer for Queues

The `QueueConnection` provides the following method for creating connection consumers for queues:

```
public ConnectionConsumer createConnectionConsumer(  
    Queue queue,  
    String messageSelector,  
    ServerSessionPool sessionPool,  
    int maxMessages  
) throws JMSEException
```

You must specify the name of the associated queue, the message selector for filtering messages, the associated server session pool for accessing server sessions, and the maximum number of messages that can be assigned to the server session simultaneously. For information about message selectors, see “[Filtering Messages](#)” on page 4-66.

For more information about the `QueueConnection` class methods, see the `javax.jms.QueueConnection` Javadoc. For more information about the `ConnectionConsumer` class, see the `javax.jms.ConnectionConsumer` Javadoc.

### Create a Connection Consumer for Topics

The `TopicConnection` provides the following two methods for creating `ConnectionConsumers` for topics:

```
public ConnectionConsumer createConnectionConsumer(  
    Topic topic,  
    String messageSelector,  
    ServerSessionPool sessionPool,  
    int maxMessages  
) throws JMSEException
```

```
public ConnectionConsumer createDurableConnectionConsumer (
    Topic topic,
    String messageSelector,
    ServerSessionPool sessionPool,
    int maxMessages
) throws JMSException
```

For each method, you must specify the name of the associated topic, the message selector for filtering messages, the associated server session pool for accessing server sessions, and the maximum number of messages that can be assigned to the server session simultaneously. For information about message selectors, see [“Filtering Messages” on page 4-66](#).

Each method creates a connection consumer; but, the second method also creates a durable connection consumer for use with durable subscribers. For more information about durable subscribers, see [“Setting Up Durable Subscriptions” on page 4-53](#).

For more information about the `TopicConnection` class methods, see the [`javax.jms.TopicConnection` Javadoc](#). For more information about the `ConnectionConsumer` class, see the [`javax.jms.ConnectionConsumer` Javadoc](#).

## Example: Setting Up a PTP Client Server Session Pool

The following example illustrates how to set up a server session pool for a JMS client. The `startup()` method is similar to the `init()` method in the `examples.jms.queue.QueueSend` example, as described in [“Example: Setting Up a PTP Application” on page 4-16](#). This method also sets up the server session pool.

The following illustrates the `startup()` method, with comments highlighting each setup step.

Include the following package on the import list to implement a server session pool application:

```
import weblogic.jms.ServerSessionPoolFactory
```

Define the session pool factory static variable required for the creation of the session pool.

```
private final static String SESSION_POOL_FACTORY=
    "weblogic.jms.ServerSessionPoolFactory:examplesJMSServer";

private QueueConnectionFactory qconFactory;
```

```
private QueueConnection qcon;
private QueueSession qsession;
private QueueSender qsender;
private Queue queue;
private ServerSessionPoolFactory sessionPoolFactory;
private ServerSessionPool sessionPool;
private ConnectionConsumer consumer;
```

Create the required JMS objects.

```
public String startup(
    String name,
    Hashtable args
) throws Exception
{
    String connectionFactory = (String)args.get("connectionFactory");
    String queueName = (String)args.get("queue");
    if (connectionFactory == null || queueName == null) {
        throw new
IllegalArgumentException("connectionFactory="+connectionFactory+
                            ", queueName="+queueName);
    }
    Context ctx = new InitialContext();
    qconFactory = (QueueConnectionFactory)
        ctx.lookup(connectionFactory);
    qcon = qconFactory.createQueueConnection();
    qsession = qcon.createQueueSession(false,
        Session.AUTO_ACKNOWLEDGE);
    queue = (Queue) ctx.lookup(queueName);
    qcon.start();
}
```

**Step 1** Look up the server session pool factory in JNDI.

```
sessionPoolFactory = (ServerSessionPoolFactory)
    ctx.lookup(SESSION_POOL_FACTORY);
```

**Step 2** Create a server session pool using the server session pool factory, as follows:

```
sessionPool = sessionPoolFactory.getServerSessionPool(qcon, 5,
    false, Session.AUTO_ACKNOWLEDGE,
    examples.jms.startup.MsgListener);
```

The code defines the following:

- `qcon` as the queue connection associated with the server session pool
- `5` as the maximum number of concurrent sessions that can be retrieved by the connection consumer (to be created in step 3)

- Sessions will be non-transacted (`false`)
- `AUTO_ACKNOWLEDGE` as the acknowledge mode
- The `examples.jms.startup.MsgListener` will be used as the message listener that is instantiated and used to receive and process messages concurrently.

Step 3 Create a connection consumer, as follows:

```
consumer = qcon.createConnectionConsumer(queue, "TRUE",
    sessionPool, 10);
```

The code defines the following:

- `queue` as the associated queue
- `TRUE` as the message selector for filtering messages
- `sessionPool` as the associated server session pool for accessing server sessions
- `10` as the maximum number of messages that can be assigned to the server session simultaneously

For more information about the JMS classes used in this example, see [“WebLogic JMS Classes” on page 2-5](#) or the `javax.jms` Javadoc.

## Example: Setting Up a Pub/Sub Client Server Session Pool

The following example illustrates how to set up a server session pool for a JMS client. The `startup()` method is similar to the `init()` method in the `examples.jms.topic.TopicSend` example, as described in [“Example: Setting Up a Pub/Sub Application” on page 4-19](#). It also sets up the server session pool.

The following illustrates `startup()` method, with comments highlighting each setup step.

Include the following package on the import list to implement a server session pool application:

```
import weblogic.jms.ServerSessionPoolFactory
```

Define the session pool factory static variable required for the creation of the session pool.

```
private final static String SESSION_POOL_FACTORY=
    "weblogic.jms.ServerSessionPoolFactory:examplesJMSServer";

private TopicConnectionFactory tconFactory;
private TopicConnection tcon;
private TopicSession tsession;
private TopicSender tsender;
private Topic topic;
private ServerSessionPoolFactory sessionPoolFactory;
private ServerSessionPool sessionPool;
private ConnectionConsumer consumer;
```

Create the required JMS objects.

```
public String startup(
    String name,
    Hashtable args
) throws Exception
{
    String connectionFactory = (String)args.get("connectionFactory");
    String topicName = (String)args.get("topic");
    if (connectionFactory == null || topicName == null) {
        throw new
IllegalArgumentException("connectionFactory="+connectionFactory+
                            ", topicName="+topicName);
    }
    Context ctx = new InitialContext();
    tconFactory = (TopicConnectionFactory)
        ctx.lookup(connectionFactory);
    tcon = tconFactory.createTopicConnection();
    tsession = tcon.createTopicSession(false,
        Session.AUTO_ACKNOWLEDGE);
    topic = (Topic) ctx.lookup(topicName);
    tcon.start();
}
```

**Step 1** Look up the server session pool factory in JNDI.

```
sessionPoolFactory = (ServerSessionPoolFactory)
    ctx.lookup(SESSION_POOL_FACTORY);
```

**Step 2** Create a server session pool using the server session pool factory, as follows:

```
sessionPool = sessionPoolFactory.getServerSessionPool(tcon, 5,
    false, Session.AUTO_ACKNOWLEDGE,
    examples.jms.startup.MsgListener);
```

The code defines the following:

- `tcon` as the topic connection associated with the server session pool
- `5` as the maximum number of concurrent sessions that can be retrieved by the connection consumer (to be created in step 3)
- Sessions will be non-transacted (`false`)
- `AUTO_ACKNOWLEDGE` as the acknowledge mode
- The `examples.jms.startup.MsgListener` will be used as the message listener that is instantiated and used to receive and process messages concurrently.

Step 3 Create a connection consumer, as follows:

```
consumer = tcon.createConnectionConsumer(topic, "TRUE",  
    sessionPool, 10);
```

The code defines the following:

- `topic` as the associated topic
- `TRUE` as the message selector for filtering messages
- `sessionPool` as the associated server session pool for accessing server sessions
- `10` as the maximum number of messages that can be assigned to the server session simultaneously

For more information about the JMS classes used in this example, see [“WebLogic JMS Classes” on page 2-5](#) or the `javax.jms` Javadoc.

## Using Multicasting

Multicasting enables the delivery of messages to a select group of hosts that subsequently forward the messages to subscribers.

The benefits of multicasting include:

- Near real-time delivery of messages to host group.

- High scalability due to the reduction in the amount of resources required by the JMS server to deliver messages to subscribers.

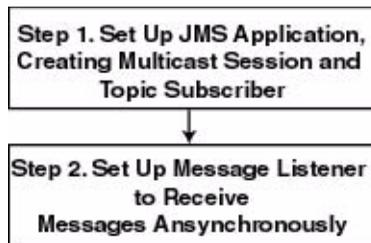
The limitations of multicasting include:

- Multicast messages are not guaranteed to be delivered to all members of the host group. For messages requiring reliable delivery and recovery, you should not use multicasting.
- For interoperability with different versions of WebLogic Server, clients cannot have an earlier release of WebLogic Server installed than the host. They must all have at least the same version or higher.

For an example of when multicasting might be useful, consider a stock ticker. When accessing stock quotes, timely delivery is more important than reliability. When accessing the stock information in real-time, if all or a portion of the contents is not delivered, the client can simply request the information to be resent. Clients would not want to have the information recovered, in this case, as by the time it is redelivered, it would be out-of-date.

The following figure illustrates the steps required to set up multicasting.

**Figure 4-5 Setting Up Multicasting**



**Note:** Multicasting is only supported for the Pub/Sub messaging model, and only for non-durable subscribers.

Monitoring statistics are not provided for multicast sessions or consumers.

Before setting up multicasting, the connection factory and destination must be configured to support multicasting, as follows:

- For each connection factory, the system administrator configures the maximum number of outstanding messages that can exist on a multicast session and whether the most recent or oldest messages are discarded in the event the

maximum is reached. If the message maximum is reached, a `DataOverrunException` is thrown, and messages are automatically discarded. These attributes are also dynamically configurable, as described in [“Dynamically Configuring Multicasting Configuration Attributes”](#) on page 4-85.

- For each destination, the Multicast Address (IP), Port, and TTL (Time-To-Live) attributes are specified. To better understand the TTL attribute setting, see [“Example: Multicast TTL”](#) on page 4-86.

**Note:** It is strongly recommended that you seek the advice of your network administrator when configuring the multicast IP address, port, and time-to-live attributes to ensure that the appropriate values are set.

For more information on the multicasting configuration attributes, see the [Administration Console Online Help](#). The multicast configuration attributes are also summarized in [Appendix A, “Configuration Checklists.”](#)

## Step 1: Set Up the JMS Application, Creating Multicast Session and Topic Subscriber

Set up the JMS application as described in [“Setting Up a JMS Application”](#) on page 4-4. However, when creating sessions, as described in [“Step 3: Create a Session Using the Connection”](#) on page 4-8, specify that the session would like to receive multicast messages by setting the `acknowledgeMode` value to `MULTICAST_NO_ACKNOWLEDGE`.

**Note:** Multicasting is only supported for the Pub/Sub messaging model for non-durable subscribers. An attempt to create a durable subscriber on a multicast session will cause a `JMSEException` to be thrown.

For example, the following method illustrates how to create a multicast session for the Pub/Sub messaging model.

```
tsession = tcon.createTopicSession(  
    false,  
    WLSession.MULTICAST_NO_ACKNOWLEDGE  
);
```

**Note:** On the client side, each multicasting session requires one dedicated thread to retrieve messages off the socket. Therefore, you should increase the JMS client-side thread pool size to adjust for this. For more information on

adjusting the thread pool size, see the “Tuning Thread Pools and EJB Pools” section in the “[WebLogic JMS Performance Guide](#)” white paper, at <http://dev2dev.bea.com/products/wlserver/resources.jsp>, which discusses tuning JMS client-side thread pools.

In addition, create a topic subscriber, as described in “[Create TopicPublishers and TopicSubscribers](#)” on page 4-12.

For example, the following code illustrates how to create a topic subscriber:

```
tsubscriber = tsession.createSubscriber(myTopic);
```

**Note:** The `createSubscriber()` method fails if the specified destination is not configured to support multicasting.

## Step 2: Set Up the Message Listener

Multicast topic subscribers can only receive messages asynchronously. If you attempt to receive synchronous messages on a multicast session, a `JMSException` is thrown.

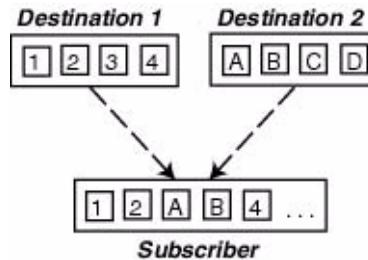
Set up the message listener for the topic subscriber, as described in “[Receiving Messages Asynchronously](#)” on page 4-30.

For example, the following code illustrates how to establish a message listener:

```
tsubscriber.setMessageListener(this);
```

When receiving messages, WebLogic JMS tracks the order in which messages are sent by the destinations. If a multicast subscriber’s message listener receives the messages out of sequence, resulting in one or more messages being skipped, a `SequenceGapException` will be delivered to the `ExceptionListener` for the session(s) present. If a skipped message is subsequently delivered, it will be discarded. For example, in the following figure, the subscriber is receiving messages from two destinations simultaneously.

Figure 4-6 Multicasting Sequence Gap



Upon receiving the “4” message from Destination 1, a `SequenceGapException` is thrown to notify the application that a message was received out of sequence. If subsequently received, the “3” message will be discarded.

**Note:** The larger the messages being exchanged, the greater the risk of encountering a `SequenceGapException`.

## Dynamically Configuring Multicasting Configuration Attributes

During configuration, for each connection factory the system administrator configures the following information to support multicasting:

- Messages maximum specifying the maximum number of outstanding messages that can exist on a multicast session.
- Overrun policy specifying whether recent or older messages are discarded in the event the messages maximum is reached.

If the messages maximum is reached, a `DataOverrunException` is thrown and messages are automatically discarded based on the overrun policy. Alternatively, you can set the messages maximum and overrun policy using the `Session` set methods.

## 4 Developing a WebLogic JMS Application

---

The following table lists the `Session` set and get methods for each dynamically configurable attribute.

**Table 4-6 Message Producer Set and Get Methods**

Attribute	Set Method	Get Method
Messages Maximum	<pre>public void setMessagesMaximum(     int messagesMaximum ) throws JMSEException</pre>	<pre>public int getMessagesMaximum( ) throws JMSEException</pre>
Overrun Policy	<pre>public void setOverrunPolicy (     int overrunPolicy ) throws JMSEException</pre>	<pre>public int getOverrunPolicy( ) throws JMSEException</pre>

**Note:** The values set using the set methods take precedence over the configured values.

For more information about these `Session` class methods, see the [weblogic.jms.extensions.WLSession](#) Javadoc. For more information on these multicast configuration attributes, see “[JMS Destinations](#)” in the Administration Console Online Help.

## Example: Multicast TTL

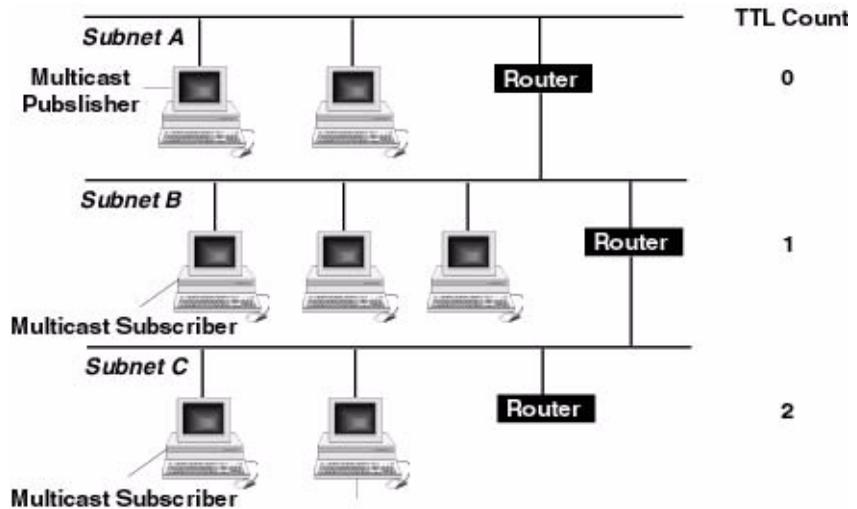
**Note:** The following example is a very simplified illustration of how the Multicast TTL (time-to-live) destination configuration attribute impacts the delivery of messages across routers. It is strongly advised that you seek the assistance of your network administrator when configuring the multicast TTL attribute to ensure that the appropriate value is set.

The Multicast TTL is independent of the message time-to-live.

The following example illustrates how the Multicast TTL destination configuration attribute impacts the delivery of messages across routers. For more information on the multicast configuration attributes, see “[JMS Destinations](#)” in the Administration Console Online Help.

Consider the following network diagram.

Figure 4-7 Multicast TTL Example



In the figure, the network consists of three subnets: Subnet A containing the multicast publisher, and Subnets B and C each containing one multicast subscriber.

If the Multicast TTL attribute is set to 0 (indicating that the messages cannot traverse any routers and are delivered on the current subnet only), when the multicast publisher on Subnet A publishes a message, the message will not be delivered to any of the multicast subscribers.

If the Multicast TTL attribute is set to 1 (indicating that messages can traverse one router), when the multicast publisher on Subnet A publishes a message, the multicast subscriber on Subnet B will receive the message.

Similarly, if the Multicast TTL attribute is set to 2 (indicating that messages can traverse two routers), when the multicast publisher on Subnet A publishes a message, the multicast subscribers on Subnets B and C will receive the message.

# Using Distributed Destinations

By enabling you to configure multiple physical destinations (queues and topics) as members of a single distributed destination set, WebLogic JMS supports service continuity in the event of a WebLogic Server failure within a cluster. Once properly configured, your producers and consumers are able to send and receive messages through the distributed destination. WebLogic JMS then distributes the messaging load across all available destination members within the distributed destination. When a member becomes unavailable due a server failure, traffic is then redirected toward other available destination members in the set.

For instructions on configuring distributed destinations using the Administration Console, see “[Configuring Distributed Destinations](#)” in the *Administration Guide*.

The following sections explain how to use distributed destinations with your JMS applications:

- [Accessing Distributed Destinations](#)
- [Accessing Distributed Destination Members](#)
- [Load Balancing Messages Across a Distributed Destination](#)
- [Distributed Destination Migration](#)

## Accessing Distributed Destinations

A distributed destination is actually a set of physical JMS destination members (queues or topics) that is accessed through a single JNDI name. As such, a distributed destination can be looked up using JNDI. It implements the `javax.jms.Destination` interface, and can be used to create producers, consumers, and browsers.

Because a distributed destination can be served by multiple WebLogic Servers within a cluster, when creating a reference to a distributed destination by using one of the `createQueue()` or `createTopic()` methods, the name supplied is simply the name of the `JMSDistributedQueueMBean` or `JMSDistributedTopicMBean` configuration MBean name. No JMS server name or separating forward slash (/) is required.

For example, the following code illustrates how to look up a distributed destination topic member:

```
topic = myTopicSession.createTopic("myDistributedTopic");
```

**Note:** When calling the `createQueue()` or `createTopic()` methods, any string containing a forward slash (/), is assumed to be the name of a distributed destination member—not a destination. If no such destination member exists, then the call will fail with an `InvalidDestinationException`.

### Looking Up Distributed Queues

A distributed queue is a set of physical JMS queue members. As such, a distributed queue can be used to create a `QueueSender`, `QueueReceiver`, and a `QueueBrowser`. The fact that a distributed queue represents multiple physical queues is mostly transparent to your application.

The queue members can be located anywhere, but must all be served by JMS servers in a single server cluster. When a message is sent to a distributed queue, it is sent to exactly one of the physical queues in the set of members for the distributed queue. Once the message arrives at the queue member, it is available for receipt by consumers of that queue member only.

**Note:** Queue members can forward messages to other queue members by configuring the Forward Delay attribute in the Administration Console, which is disabled by default. This attribute defines the amount of time, in seconds, that a distributed queue member with messages, but which has no consumers, will wait before forwarding its messages to other queue members that do have consumers.

### QueueSenders

After creating a queue sender, if the queue supplied at creation time was a distributed queue, then each time a message is produced using the sender a decision is made as to which queue member will receive the message. Each message is sent to a single physical queue member.

The message is not replicated in any way. As such, the message is only available from the queue member where it was sent. If that physical queue becomes unavailable before a given message is received, then the message is unavailable until that queue member comes back online.

It is not enough to send a message to a distributed queue and expect the message to be received by a queue receiver of that distributed queue. Since the message is sent to only one physical queue member, there must be a queue receiver receiving or listening on that queue member.

**Note:** For information on the load-balancing heuristics for distributed queues with zero consumers, see [“Load Balancing Heuristics” on page 4-96](#).

### QueueReceivers

When creating a queue receiver, if the supplied queue is a distributed queue, then a single physical queue member is chosen for the receiver at creation time. The created `QueueReceiver` is pinned to that queue member until the queue receiver loses its access to the queue member. At that point, the consumer will receive a `JMSEException`, as follows:

- If the queue receiver is synchronous, then the exception is returned to the user directly.
- If the queue receiver is asynchronous, then the exception is delivered inside of a `ConsumerClosedException` that is delivered to the `ExceptionListener` defined for the consumer session, if any.

Upon receiving such an exception, an application can close its queue receiver and recreate it. If any other queue members are available within the distributed queue, then the creation will succeed and the new queue receiver will be pinned to one of those queue members. If no other queue member is available, then the application won't be able to recreate the queue receiver and will have to try again later.

**Note:** For information on the load-balancing heuristics for distributed queues with zero consumers, see [“Load Balancing Heuristics” on page 4-96](#).

### QueueBrowsers

When creating a queue browser, if the supplied queue is a distributed queue, then a single physical queue member is chosen for the browser at creation time. The created queue browser is pinned to that queue member until the receiver loses its access to the queue member. At that point, any calls to the queue browser will receive a `JMSEException`. Any calls to the enumeration will return a `NoSuchElementException`.

**Note:** The queue browser can only browse the queue member that it is pinned to. Even though a distributed queue was specified at creation time, the queue browser cannot see or browse messages for the other queue members in the distributed destination, even if Server Affinity is enabled for the distributed destination.

## Looking Up Distributed Topics

A distributed topic is a set of physical JMS topic members. As such, a distributed topic can be used to create a `TopicPublisher` and `TopicSubscriber`. The fact that a distributed topic represents multiple physical topics is mostly transparent to the application.

**Note:** Durable subscribers (`DurableTopicSubscriber`) cannot be created for distributed topics. However, you can still create a durable subscription on distributed topic member and the other topic members will forward the messages to the topic member that has the durable subscription.

The topic members can be located anywhere but must all be served either by a single WebLogic Server or any number of servers in a cluster. When a message is sent to a distributed topic, it is sent to all of the topic members in the distributed topic set. This allows all subscribers to the distributed topic to receive messages published for the distributed topic.

A message published directly to a topic member of a distributed destination (that is, the publisher did not specify the distributed destination) is also forwarded to all the members of that distributed topic. This includes subscribers that originally subscribed to the distributed topic, and which happened to be assigned to that particular topic member. In other words, publishing a message to a specific distributed topic member automatically forwards it to all the other distributed topic members, just as publishing a message to a distributed topic automatically forwards it to all of its distributed topic members. For more information about looking up specific distributed destination members, see [“Accessing Distributed Destination Members” on page 4-94](#).

## Deploying a Message-Driven Bean on a Distributed Topic

When a MDB is deployed on a distributed topic and is targeted to a WebLogic Server instance in a cluster that is hosting two members of the distributed topic on a JMS server, the MDB gets deployed on both the members of the distributed topic. This occurs because MDBs are pinned to a distributed topic member’s destination name.

Therefore, you will receive  $[number\ of\ messages\ sent] * [number\ of\ distributed\ topic\ members]$  more messages per MDB, depending on how many distributed topic members are deployed on a WebLogic Server instance. For example, if a JMS server contains two distributed topic members, then two MDBs are deployed, one for each member, so you will receive twice as many messages.

### TopicPublishers

When creating a topic publisher, if the supplied destination is a distributed destination, then any messages sent to that distributed destination are sent to all available topic members for that distributed topic, as follows:

- If one or more of the distributed topic members is not reachable, and the message being sent is non-persistent, then the message is sent only to the available topic members.
- If one or more of the distributed topic members is not reachable, and the message being sent is persistent, then the message is stored and forwarded to the other topic members when they become reachable. However, the message can only be persistently stored if the topic member has a JMS store configured.  
**Note:** Every effort is made to first forward the message to distributed members that utilize a persistent store. However, if none of the distributed members utilize a store, then the message is still sent to one of the members according to the selected load-balancing algorithm, as described in [“Load Balancing Messages Across a Distributed Destination”](#) on page 4-94.
- If all of the distributed topic members are unreachable (regardless of whether the message is persistent or non-persistent), then the publisher receives a `JMSException` when it tries to send a message.

### TopicSubscribers

When creating a topic subscriber, if the supplied topic is a distributed topic, then the topic subscriber receives messages published to that distributed topic. If one or more of the topic members for the distributed topic are not reachable by a topic subscriber, then depending on whether the messages are persistent or non-persistent the following occurs:

- Any persistent messages published to one or more unreachable distributed topic members are eventually received by topic subscribers of those topic members

once they become reachable. However, the messages can only be persistently stored if the topic member has a JMS store configured.

- Any non-persistent messages published to those unreachable distributed topic members will not be received by that topic subscriber.

**Note:** If a JMS store is configured for a JMS server that is hosting a distributed topic member, then all the Distributed Topic System Subscribers associated with that member destination are treated as durable subscriptions, even when a topic member does not have a JMS store explicitly configured. As such, the saving of all the messages sent to these distributed topic subscribers in memory can result in unexpected memory and disk consumption. Therefore, a recommended best design practice when deploying distributed destination is to consistently configure all member destinations: either with a JMS store for durable messages, or without a JMS store for non-durable messages. For example, if you want all of your distributed topic subscribers to be non-durable, but some member destinations implicitly have a JMS store configured because their associated JMS server uses a JMS store, then you need to explicitly set the `StoreEnabled` attribute to `False` for each member destination to override the JMS server setting.

Ultimately, a topic subscriber is pinned to a physical topic member. If that topic member becomes unavailable, then the topic subscriber will receive a `JMSException`, as follows:

- If the topic subscriber is synchronous, then the exception is returned to the user directly.
- If the topic subscriber is asynchronous, then the exception is delivered inside of a `ConsumerClosedException` that is delivered to the `ExceptionListener` defined for the consumer session, if any.

Upon receiving such an exception, an application can close its topic subscriber and recreate it. If any other topic member is available within the distributed topic, then the creation should be successful and the new topic subscriber will be pinned to one of those topic members. If no other topic member is available, then the application will not be able to recreate the topic subscriber and will have to try again later.

# Accessing Distributed Destination Members

In order to access a destination member within a distributed destination, you must look up the destination member using the configured JNDI name, or supply the JMS server name and the `JMSQueueMBean` or `JMSTopicMBean` configuration MBean name, separated by a forward slash (/), to one of the `createQueue()` or `createTopic()` methods.

For example, the following code illustrates how to look up a particular member of a distributed queue (`myQueue`), on a JMS server (`myServer`):

```
queue = myQueueSession.createQueue("myServer/myQueue");
```

**Note:** When calling the `createQueue()` or `createTopic()` methods, any string containing a forward slash (/), is assumed to be the name of a distributed destination member—not a destination. If no such destination member exists, then the call will fail with an `InvalidDestinationException`.

# Load Balancing Messages Across a Distributed Destination

By using distributed destinations, WebLogic JMS can spread or balance the messaging load across multiple physical destinations, which can result in better use of resources and improved response times. The WebLogic JMS load-balancing algorithm determines the physical destinations that messages are sent to, as well as the physical destinations that consumers are assigned to.

For more information about configuring load balancing for a distributed destination, see “[Configuring Message Load Balancing](#)” in the *Administration Guide*.

## Load Balancing Options

WebLogic JMS supports two different algorithms for balancing the message load across multiple physical destinations within a given distributed destination set. You can select one of these load balancing options configuring a distributed topic or queue on the Administration Console.

- [Round-Robin Distribution](#)

- [Random Distribution](#)

## Round-Robin Distribution

In the round-robin algorithm, WebLogic JMS maintains an ordering of physical destinations within the distributed destination. The messaging load is distributed across the physical destinations one at a time in the order that they are defined in the WebLogic Server configuration (`config.xml`) file. Each WebLogic Server maintains an identical ordering, but may be at a different point within the ordering. Multiple threads of execution within a single server using a given distributed destination affect each other with respect to which physical destination a member is assigned to each time they produce a message. Round-robin is the default algorithm and doesn't need to be configured.

If weights are assigned to any of the physical destinations in the set for a given distributed destination, then those physical destinations appear multiple times in the ordering. For instance, if the weights of destinations A, B and C are 2, 5, and 3 respectively, then the ordering will be A, B, C, A, B, C, B, C, B, B. That is, a number of passes are made through the basic ordering (A, B, C). The number of passes is equal to the highest weight of the destinations within the set. On each pass, only those destinations with a weight that is greater than or equal to the ordinal value of the pass are included in the ordering. Following this logic, this example would produce the following results:

- A is dropped from the ordering after two passes.
- C is dropped after three passes.
- B is the only one remaining on the fourth and fifth passes.

## Random Distribution

The random distribution algorithm uses the weight assigned to the physical destinations to compute a weighted distribution for the set of physical destinations. The messaging load is distributed across the physical destinations by pseudo-randomly accessing the distribution. In the short run, the load will not be directly proportional to the weight. In the long run, the distribution will approach the limit of the distribution. A pure random distribution can be achieved by setting all the weights to the same value, which is typically 1.

Adding or removing a member (either administratively or as a result of a WebLogic Server shutdown/restart event) requires a recomputation of the distribution. Such events should be infrequent however, and the computation is generally simple, running in  $O(n)$  time.

### Consumer Load Balancing

When an application creates a consumer, it must provide a destination. If that destination represents a distributed destination, then WebLogic JMS must find a physical destination that consumer will receive messages from. The choice of which destination member to use is made by using one of the load-balancing algorithms described in [“Load Balancing Options” on page 4-94](#). The choice is made only once: when the consumer is created. From that point on, the consumer gets messages from that member only.

### Producer Load Balancing

When a producer sends a message, WebLogic JMS looks at the destination where the message is being sent. If the destination is a distributed destination, WebLogic JMS makes a decision as to where the message will be sent. That is, the producer will send to one of the destination members according to one of the load-balancing algorithms described in [“Load Balancing Options” on page 4-94](#).

The producer makes such a decision each time it sends a message. However, there is no compromise of ordering guarantees between a consumer and producer, because consumers are load balanced once, and are then pinned to a single destination member.

**Note:** If a producer attempts to send a persistent message to a distributed destination, every effort is made to first forward the message to distributed members that utilize a persistent store. However, if none of the distributed members utilize a persistent store, then the message will still be sent to one of the members according to the selected load-balancing algorithm.

### Load Balancing Heuristics

In addition to the algorithms described in [“Load Balancing Options” on page 4-94](#), WebLogic JMS uses the following heuristics when choosing an instance of a destination.

- [Transaction Affinity](#)

- [Server Affinity](#)
- [Queues with Zero Consumers](#)

## Transaction Affinity

When producing multiple messages within a transacted session, an effort is made to send all messages produced to the same WebLogic Server. Specifically, if a session sends multiple messages to a single distributed destination, then all of the messages are routed to the same physical destination. If a session sends multiple messages to multiple different distributed destinations, an effort is made to choose a set of physical destinations served by the same WebLogic Server.

## Server Affinity

When the Server Affinity option is enabled for distributed destinations, then before a WebLogic Server instance attempts to load balance consumers or producers across all the members of a distributed destination in a domain, it will first attempt to load balance across any local members that are running on the same WebLogic Server instance.

For more information about configuring server affinity for a distributed destination using the Server Affinity Enabled option on the JMS connection factory, see “[Configuring Server Affinity](#)” in the *Administration Guide*.

**Note:** The Server Affinity Enabled attribute does not affect queue browsers. Therefore, a queue browser created on a distributed queue can be pinned to a remote distributed queue member even when Server Affinity is enabled.

## Queues with Zero Consumers

When load balancing consumers across multiple remote physical queues, if one or more of the queues have zero consumers, then those queues alone are considered for balancing the load. Once all the physical queues in the set have at least one consumer, the standard algorithms apply.

In addition, when producers are sending messages, queues with zero consumers are not considered for message production, unless all instances of the given queue have zero consumers.

### Defeating Load Balancing

Applications can defeat load balancing by directly accessing the individual physical destinations. That is, if the physical destination has no JNDI name, it can still be referenced using the `createQueue()` or `createTopic()` methods.

- [JNDI Lookup](#)
- [CreateQueue\(\) and CreateTopic\(\)](#)
- [Connection Factories](#)

#### JNDI Lookup

If a physical destination has a JNDI name, then it can be looked up using JNDI. The returned destination can then be used to create a consumer or receiver.

#### createQueue() and createTopic()

An application can also obtain a reference to a topic or queue using the `createQueue()` and `createTopic()` methods. When using these methods, the application must supply a vendor-specific string identifying the destination that they want a reference to. The vendor-specific string for WebLogic JMS is of the form *server/destination*, where “server” is the name of a JMS server and “destination” is the name of a queue or topic on that JMS server.

#### Connection Factories

Applications that use distributed destinations to distribute or balance their producers and consumers across multiple physical destinations, but do not want to make a load balancing decision each time a message is produced, can use a connection factory with the Load Balancing Enabled attribute disabled (i.e., set to False).

For more information about configuring load balancing for a distributed destination, see “[Enabling Message Load Balancing](#)” in the *Administration Guide*.

## How Distributed Destination Load Balancing Is Affected When Using the “Server Affinity Enabled” Attribute

The following table explains how the setting of the JMS connection factory’s Server Affinity Enabled attribute affects the load balancing preferences for distributed destination members. The order of preference depends on the type of operation and whether or not durable subscriptions or persistent messages are involved.

The Server Affinity Enabled attribute for distributed destinations is different from the server affinity provided by the Default Load Algorithm attribute in the `ClusterMBean`, which is also used by the JMS connection factory to create initial context affinity for client connections. For more information, refer to “[Load Balancing for EJBs and RMI Objects](#)” in *Using WebLogic Server Clusters*.

**Table 4-7 Server Affinity Load Balancing Preferences**

When the operation is...	and Server Affinity Enabled is...	then load balancing preference is given to a...
<ul style="list-style-type: none"> <li>■ <code>createReceiver()</code> for queues</li> <li>■ <code>createDurableSubscriber()</code> for topics</li> </ul>	True	<ol style="list-style-type: none"> <li>1. local member without a consumer</li> <li>2. local member</li> <li>3. remote member without a consumer</li> <li>4. remote member</li> </ol>
<code>createReceiver()</code> for queues	False	<ol style="list-style-type: none"> <li>1. member without a consumer</li> <li>2. member</li> </ol>
<code>createSubscriber()</code> for topics ( <b>Note:</b> non-durable subscribers)	True or False	<ol style="list-style-type: none"> <li>1. local member without a consumer</li> <li>2. local member</li> <li>3. local proxy topic (see “<a href="#">Monitoring Distributed Destination System Subscriptions and Proxy Topic Members</a>” in the <i>Administration Guide</i>)</li> </ol>

## 4 Developing a WebLogic JMS Application

**Table 4-7 Server Affinity Load Balancing Preferences**

When the operation is...	and Server Affinity Enabled is...	then load balancing preference is given to a...
<ul style="list-style-type: none"> <li>■ <code>createSender()</code> for queues</li> <li>■ <code>createPublisher()</code> for topics</li> </ul>	True or False	<p>There is no separate machinery for load balancing a JMS producer creation. JMS producers are created on the server where your JMS connection is load balanced or pinned to.</p> <p>For more information about load balancing JMS connections created via a connection factory, refer to the “<a href="#">Load Balancing for EJBs and RMI Objects</a>” in <i>Using WebLogic Server Clusters</i>.</p>
<p>For persistent messages:</p> <ul style="list-style-type: none"> <li>■ <code>QueueSender.send()</code></li> </ul>	True	<ol style="list-style-type: none"> <li>1. local member with a consumer and a store</li> <li>2. remote member with a consumer and a store</li> <li>3. local member with a store</li> <li>4. remote member with a store</li> <li>5. local member with a consumer</li> <li>6. remote member with a consumer</li> <li>7. local member</li> <li>8. remote member</li> </ol>
<p>For persistent messages:</p> <ul style="list-style-type: none"> <li>■ <code>QueueSender.send()</code></li> </ul>	False	<ol style="list-style-type: none"> <li>1. member with a consumer and a store</li> <li>2. member with a store</li> <li>3. member with a consumer</li> <li>4. member</li> </ol>
<p>For non-persistent messages:</p> <ul style="list-style-type: none"> <li>■ <code>QueueSender.send()</code></li> </ul>	True	<ol style="list-style-type: none"> <li>1. local member with a consumer</li> <li>2. remote member with a consumer</li> <li>3. local member</li> <li>4. remote member</li> </ol>

Table 4-7 Server Affinity Load Balancing Preferences

When the operation is...	and Server Affinity Enabled is...	then load balancing preference is given to a...
For non-persistent messages: ■ <code>QueueSender.send()</code>	False	<ol style="list-style-type: none"> <li>1. member with a consumer</li> <li>2. member</li> </ol>
<code>createConnectionConsumer()</code> for session pool queues and topics	True or False	<ol style="list-style-type: none"> <li>1. local member <i>only</i></li> </ol> <p><b>Note:</b> Session pools are now used rarely, as they are not a required part of the J2EE specification, do not support JTA user transactions, and are largely superseded by message-driven beans (MDBs), which are simpler, easier to manage, and more capable.</p>

## Distributed Destination Migration

For JMS implementations that take use the WebLogic Server 7.0 service migration feature, when a JMS server fails it can migrate to another WebLogic Server within a cluster—along with all of its distributed destination members. However, the target WebLogic Server may already be hosting a JMS server with all of its physical destinations. This can lead to situations where the same WebLogic Server hosts two physical destinations for a single distributed destination. This is permissible in the short term, since a WebLogic Server can host multiple physical destinations for that distributed destination. However, load balancing in this situation is less effective.

In such a situation, each JMS server on a target WebLogic Server operates independently. This is necessary to avoid merging of the two destination instances, and/or disabling of one instance, which can make some messages unavailable for a prolonged period of time. The long-term intent, however, is to eventually re-migrate the migrated JMS server to yet another WebLogic Server in the cluster.

For more information about the configuring JMS migratable targets, see [“Configuring JMS Migratable Targets” on page 3-6](#).

### **Distributed Destination Failover**

If the server instance that is hosting the JMS connections for the JMS producers and JMS consumers should fail, then all the producers and consumers using these connections are closed and are *not* re-created on another server instance in the cluster. Furthermore, if a server instance that is hosting a JMS destination should fail, then all the JMS consumers for that destination are closed and not re-created on another server instance in the cluster.

If the distributed queue member on which a queue producer is created should fail, yet the WebLogic Server instance where the producer's JMS connection resides is still running, the producer remains alive and WebLogic JMS will fail it over to another distributed queue member, irrespective of whether the Load Balancing option is enabled.

For more information about procedures for recovering from a WebLogic Server failure, see "[Managing JMS](#)" in the *Administration Guide*.

# 5 Using Transactions with WebLogic JMS

The following sections describe how to use transactions with WebLogic JMS:

- [Overview of Transactions](#)
- [Using JMS Transacted Sessions](#)
- [Using JTA User Transactions](#)
- [Asynchronous Messaging Within JTA User Transactions Using Message Driven Beans](#)
- [Example: JMS and EJB in a JTA User Transaction](#)

**Note:** For more information about the JMS classes described in this section, access the latest JMS Specification and Javadoc supplied on the Sun Microsystems' Java Web site at the following location:  
<http://java.sun.com/products/jms/docs.html>

## Overview of Transactions

A transaction enables an application to coordinate a group of messages for production and consumption, treating messages sent or received as an atomic unit.

When an application commits a transaction, all of the messages it received within the transaction are removed from the messaging system and the messages it sent within the transaction are actually delivered. If the application rolls back the transaction, the messages it received within the transaction are returned to the messaging system and messages it sent are discarded.

When a topic subscriber rolls back a received message, the message is redelivered to that subscriber. When a queue receiver rolls back a received message, the message is redelivered to the queue, not the consumer, so that another consumer on that queue may receive the message.

For example, when shopping online, you select items and store them in an online shopping cart. Each ordered item is stored as part of the transaction, but your credit card is not charged until you confirm the order by checking out. At any time, you can cancel your order and empty your cart, rolling back all orders within the current transaction.

There are three ways to use transactions with JMS:

- If you are using only JMS in your transactions, you can create a *JMS transacted session*.
- If you are mixing other operations, such as EJB, with JMS operations, you should use a *Java Transaction API (JTA) user transaction* in a non-transacted JMS session.
- Use message driven beans.

To enable multiple JMS servers in the same JTA user transaction, or to combine JMS operations with non-JMS operations (such as EJB), the two-phase commit license is required. For more information, see [“Using JTA User Transactions” on page 5-5](#).

The following sections explain how to use a JMS transacted session and JTA user transaction.

**Note:** When using transactions, it is recommended that you define a session exception listener to handle any problems that occur before a transaction is committed or rolled back, as described in [“Defining a Session Exception Listener” on page 4-48](#).

If the `acknowledge()` method is called within a transaction, it is ignored. If the `recover()` method is called within a transaction, a `JMSException` is thrown.

# Using JMS Transacted Sessions

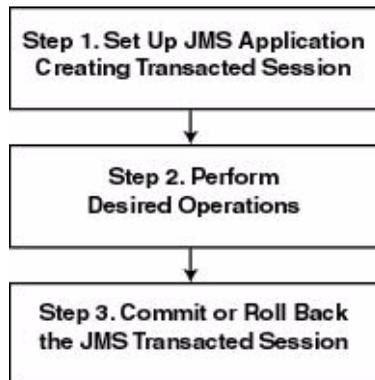
A JMS transacted session supports transactions that are located within the session. A JMS transacted session's transaction will not have any effects outside of the session. For example, rolling back a session will roll back all sends and receives on that session, but will not roll back any database updates. JTA user transactions are ignored by JMS transacted sessions.

Transactions in JMS transacted sessions are started implicitly, after the first occurrence of a send or receive operation, and chained together—whenever you commit or roll back a transaction, another transaction automatically begins.

Before using a JMS transacted session, the system administrator should adjust the connection factory (Transaction Timeout) and/or session pool (Transaction) attributes, as necessary for the application development environment, as described in [Managing JMS](#) in the *Administration Guide*.

The following figure illustrates the steps required to set up and use a JMS transacted session.

**Figure 5-1 Setting Up and Using a JMS Transacted Session**



# Step 1: Set Up JMS Application, Creating Transacted Session

Set up the JMS application as described in [“Setting Up a JMS Application” on page 4-4](#), however, when creating sessions, as described in [“Step 3: Create a Session Using the Connection” on page 4-8](#), specify that the session is to be transacted by setting the `transacted` boolean value to `true`.

For example, the following methods illustrate how to create a transacted session for the PTP and Pub/sub messaging models, respectively:

```
qsession = qcon.createQueueSession(  
    true,  
    Session.AUTO_ACKNOWLEDGE  
);  
  
tsession = tcon.createTopicSession(  
    true,  
    Session.AUTO_ACKNOWLEDGE  
);
```

Once defined, you can determine whether or not a session is transacted using the following session method:

```
public boolean getTransacted(  
    ) throws JMSEException
```

**Note:** The acknowledge value is ignored for transacted sessions.

# Step 2: Perform Desired Operations

Perform the desired operations associated with the current transaction.

# Step 3: Commit or Roll Back the JMS Transacted Session

Once you have performed the desired operations, execute one of the following methods to commit or roll back the transaction.

To commit the transaction, execute the following method:

```
public void commit(  
    ) throws JMSEException
```

The `commit()` method commits all messages sent or received during the current transaction. Sent messages are made visible, while received messages are removed from the messaging system.

To roll back the transaction, execute the following method:

```
public void rollback(  
    ) throws JMSEException
```

The `rollback()` method cancels any messages sent during the current transaction and returns any messages received to the messaging system.

If either the `commit()` or `rollback()` methods are issued outside of a JMS transacted session, a `IllegalStateException` is thrown.

## Using JTA User Transactions

The Java Transaction API (JTA) supports transactions across multiple data resources. JTA is implemented as part of WebLogic Server and provides a standard Java interface for implementing transaction management.

You program your JTA user transaction applications using the `javax.transaction.UserTransaction` object to begin, commit, and roll back the transactions. When mixing JMS and EJB within a JTA user transaction, you can also start the transaction from the EJB, as described in “[Transactions in EJB Applications](#)” in *Programming WebLogic JTA*.

You can start a JTA user transaction after a transacted session has been started; however, the JTA transaction will be ignored by the session and vice versa.

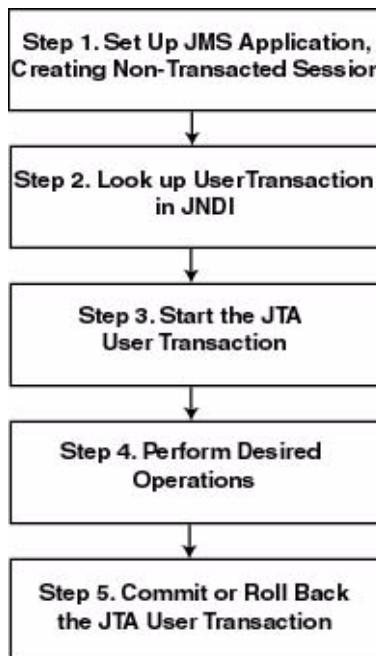
WebLogic Server supports the two-phase commit protocol (2PC), enabling an application to coordinate a single JTA transaction across two or more resource managers. It guarantees data integrity by ensuring that transactional updates are committed in all of the participating resource managers, or are fully rolled back out of all the resource managers, reverting to the state prior to the start of the transaction.

**Note:** A separate 2PC transaction license is required to support this protocol. For transaction migration considerations related to 2PC, see [“Porting WebLogic JMS Applications”](#) on page 6-1.

Before using a JTA transacted session, the system administrator must configure the connection factories to support JTA user transactions by selecting the User Transactions Enabled check box, as described in [“JMS Connection Factories”](#) in the *Administration Console Online Help*.

The following figure illustrates the steps required to set up and use a JTA user transaction.

**Figure 5-2 Setting Up and Using a JTA User Transaction**



## Step 1: Set Up JMS Application, Creating Non-Transacted Session

Set up the JMS application as described in [“Setting Up a JMS Application”](#) on page 4-4, however, when creating sessions, as described in [“Step 3: Create a Session Using the Connection”](#) on page 4-8, specify that the session is to be non-transacted by setting the `transacted` boolean value to `false`.

For example, the following methods illustrate how to create a non-transacted session for the PTP and Pub/sub messaging models, respectively.

```
qsession = qcon.createQueueSession(
    false,
    Session.AUTO_ACKNOWLEDGE
);

tsession = tcon.createTopicSession(
    false,
    Session.AUTO_ACKNOWLEDGE
);
```

**Note:** When a user transaction is active, the acknowledge mode is ignored.

## Step 2: Look Up User Transaction in JNDI

The application uses JNDI to return an object reference to the `UserTransaction` object for the WebLogic Server domain.

You can look up the `UserTransaction` object by establishing a JNDI context (`context`) and executing the following code, for example:

```
UserTransaction xact =
    ctx.lookup("javax.transaction.UserTransaction");
```

## Step 3: Start the JTA User Transaction

Start the JTA user transaction using the `UserTransaction.begin()` method. For example:

```
xact.begin();
```

### Step 4: Perform Desired Operations

Perform the desired operations associated with the current transaction.

### Step 5: Commit or Roll Back the JTA User Transaction

Once you have performed the desired operations, execute one of the following `commit()` or `rollback()` methods on the `UserTransaction` object to commit or roll back the JTA user transaction.

To commit the transaction, execute the following method:

```
xact.commit();
```

The `commit()` method causes WebLogic Server to call the Transaction Manager to complete the transaction, and commit all operations performed during the current transaction. The Transaction Manager is responsible for coordinating with the resource managers to update any databases.

To roll back the transaction, execute the following method:

```
xact.rollback();
```

The `rollback()` method causes WebLogic Server to call the Transaction Manager to cancel the transaction, and roll back all operations performed during the current transactions.

Once you call the `commit()` or `rollback()` method, you can optionally start another transaction by calling `xact.begin()`.

# Asynchronous Messaging Within JTA User Transactions Using Message Driven Beans

Because JMS cannot determine which, if any, transaction to use for an asynchronously delivered message, JMS asynchronous message delivery is not supported within JTA user transactions.

However, message driven beans provide an alternative approach. A message driven bean can automatically begin a user transaction just prior to message delivery.

For information on using message driven beans to simulate asynchronous message delivery, see “[Designing Message-Driven Beans](#)” in *Programming WebLogic EJB*.

## Example: JMS and EJB in a JTA User Transaction

The following example shows how to set up an application for mixed EJB and JMS operations in a JTA user transaction by looking up a `javax.transaction.UserTransaction` using JNDI, and beginning and then committing a JTA user transaction. In order for this example to run, the User Transactions Enabled check box must be selected when the system administrator configures the connection factory.

**Note:** In addition to this simple JTA User Transaction example, refer to the example provided with WebLogic JTA, located in the `WL_HOME\samples\server\src\examples\jta\jmsjdcdb` directory, where where `WL_HOME` is the top-level directory of your WebLogic Platform installation.

Import the appropriate packages, including the `javax.transaction.UserTransaction` package.

```
import java.io.*;
import java.util.*;
import javax.transaction.UserTransaction;
import javax.naming.*;
import javax.jms.*;
```

Define the required variables, including the JTA user transaction variable.

```
public final static String JTA_USER_XACT=
    "javax.transaction.UserTransaction";
    .
    .
    .
```

- Step 1** Set up the JMS application, creating a non-transacted session. For more information on setting up the JMS application, refer to [“Setting Up a JMS Application” on page 4-4](#).

```
//JMS application setup steps including, for example:
```

```
    qsession = qcon.createQueueSession(false,
        Session.CLIENT_ACKNOWLEDGE);
```

- Step 2** Look up the `UserTransaction` using JNDI.

```
UserTransaction xact = (UserTransaction)
    ctx.lookup(JTA_USER_XACT);
```

- Step 3** Start the JTA user transaction.

```
xact.begin();
```

- Step 4** Perform the desired operations.

```
// Perform some JMS and EJB operations here.
```

- Step 5** Commit the JTA user transaction.

```
xact.commit();
```

# 6 Porting WebLogic JMS Applications

The following sections describe how to port your WebLogic JMS applications to a newer version of WebLogic Server:

- [Existing Feature Functionality Changes](#)
- [Porting Existing Applications](#)
- [Deleting JDBC Database Stores](#)

## Existing Feature Functionality Changes

Changes in existing feature functionality have been made in order to comply with [Sun Microsystems's JMS Specification](#). Therefore, you should check feature functionality changes in the following tables before beginning any porting procedures:

- [Existing Feature 5.1 to 6.0 Functionality Changes](#)
- [Existing Feature 6.0 to 6.1 Functionality Changes](#)

## Existing Feature 5.1 to 6.0 Functionality Changes

The following table lists the changes in existing feature functionality from WebLogic Server version 5.1, and also indicates any code changes that might be required as a result. For additional information pertaining to the JMS Specification’s version change history, refer to Chapter 11, “Change History” in the specification.

Category	Description	Code Modification
Connection Factories	<p>Two default connection factories have been deprecated. The JNDI names for these factories are:</p> <ul style="list-style-type: none"> <li>■ <code>javax.jms.QueueConnectionFactory</code></li> <li>■ <code>javax.jms.TopicConnectionFactory</code></li> </ul> <p>For backwards compatibility, the JNDI names for these two connection factories are still defined and supported.</p> <p>WebLogic JMS 6.x or later defines one connection factory, by default: <code>weblogic.jms.ConnectionFactory</code></p> <p>You can also specify user-defined connection factories using the <i>Administration Console</i>.</p> <p><b>Note:</b> Using the default connection factory, you have no control over the WebLogic server on which the connection factory may be deployed. If you would like to target a particular WebLogic Server, create a new connection factory and specify the appropriate WebLogic Server target(s).</p>	<p>It is recommended that existing code that use the deprecated classes be modified to use a new default or user-defined connection factory class.</p> <p>For example, if your code specified the following constant using the default queue connection factory:</p> <pre>public final static String JMS_FACTORY="javax.jms.QueueCon nectionFactory"</pre> <p>You should modify the constant to use a new user-defined connection factory, for example:</p> <pre>public final static String JMS_FACTORY="weblogic.jms.Queue ConnectionFactory"</pre> <p>For true backwards compatibility with previous releases, you should ensure that you select the Allow Close In onMessage and User Transactions Enabled check boxes when configuring the connection factory.</p> <p>For more information about defining connection factories, see “<a href="#">JMS Connection Factories</a>” in the <i>Administration Console Online Help</i>.</p>
	<p>In order to instantiate the default connection factory on a particular WebLogic Server, you must select the Enable Default JMS Connection Factories check box when configuring the WebLogic Server.</p>	<p>None required. This is a configuration requirement. For more information, see “<a href="#">Server</a>” in the <i>Administration Console Online Help</i>.</p>

Category	Description	Code Modification
Connections	When closing a connection, the call blocks until outstanding synchronous calls and asynchronous listeners have completed.	None required.
Sessions	When closing a session, the call blocks until outstanding synchronous calls and asynchronous listeners have completed.	None required.
Message Consumers	If multiple topic subscribers are defined in the same session for the same topic, each consumer will receive its own copy of a message.	None required.
	When closing a message consumer, the call blocks until the method call completes and any outstanding synchronous applications are cancelled.	None required.
	In order to comply with the JMS specification, if the <code>close()</code> method is called from within an <code>onMessage()</code> method, the application will hang unless the Allow Close In OnMessage check box is selected when configuring the connection factory. If the acknowledge mode is <code>AUTO_ACKNOWLEDGE</code> , the current message will still be automatically acknowledged.	None required. This is a configuration requirement. For more information, see “ <a href="#">JMS Connection Factories</a> ” in the <i>Administration Console Online Help</i> .

## 6 Porting WebLogic JMS Applications

---

Category	Description	Code Modification
Message Header Field	The <code>JMSMessageID</code> header field format has changed.	<p>If you wish to access existing messages using the <code>JMSMessageID</code>, you may need to run one of the following</p> <pre>weblogic.jms.extensions.JMSHelper</pre> <p>methods to convert between WebLogic pre-JMS 5.1 and JMS 6.x <code>JMSMessageID</code> formats.</p> <p>To convert from pre-5.1 to 6.x <code>JMSMessageID</code> format:</p> <pre>public void oldJMSMessageIDToNew(     String id,     long timeStamp ) throws JMSEException</pre> <p>To convert from 6.1 to pre-6.1 <code>JMSMessageID</code> format:</p> <pre>public void newJMSMessageIDToOld(     String id,     long timeStamp ) throws JMSEException</pre>

---

Category	Description	Code Modification
Destinations	The <code>createQueue()</code> and <code>createTopic()</code> methods do not create destinations dynamically, only references to destinations that already exist given the vendor-specific destination name.	<p>Update any portion of code that uses <code>createQueue()</code> or <code>createTopic()</code> to dynamically create destinations using the following JMSHelper class methods, respectively:</p> <pre>createPermanentQueueAsync() and createPermanentTopicAsync().</pre> <p>For example, if your code used the following method to dynamically create a queue:</p> <pre>queue=qsession.createQueue(queueName);</pre> <p>You should modify the code to dynamically create a queue, as described in the sample <code>findQueue()</code> method in <a href="#">“Using the JMSHelper Class Methods”</a> on page 4-50.</p> <p>For more information on the JMSHelper classes, see <a href="#">“Creating Destinations Dynamically”</a> on page 4-49.</p>
	When creating temporary destinations, you must specify a temporary template.	None required. This is a configuration requirement. For more information, see <a href="#">“JMS Templates”</a> in the <i>Administration Console</i> Online Help.
	You must be the owner of the connection in order to create a message consumer for that temporary destination.	When creating a message consumer on a temporary destination, ensure that you are the owner of the connection.
Durable Subscribers	You no longer need to manually create JDBC tables for durable subscribers. They are created automatically.	None required.
	There is no limit on the number of durable subscribers that can be created.	None required.
	When defining a client ID programatically, it must be defined <i>immediately</i> after creating a connection. Otherwise, an exception will be thrown and you will be unable to make any other JMS calls on that connection.	Ensure that the <code>setClientID()</code> method is issued immediately after creating the connection. For more information, refer to <a href="#">“Defining the Client ID”</a> on page 4-55.

## 6 Porting WebLogic JMS Applications

---

Category	Description	Code Modification
Session Pools	Session pool factories, session pools, referenced connection factories, referenced destinations, and associated connection consumers must all be targeted on the same JMS server.	Ensure that all objects are targeted on the same JMS server.
	The <code>SessionPoolManager</code> and <code>ConnectionConsumerManager</code> interfaces that were published as part of the WebLogic JMS version 5.1 Javadoc have been removed from the version 6.x and later Javadoc, as they are system interfaces and should not be used within client applications.	If used, remove any references to these objects from the client application.
Transactions	To combine JMS and EJB database calls within the same transaction, a two-phase commit (2PC) license is required. In previous releases of WebLogic Server, it was possible to combine them by using the same database connection pool.	None required.
	Recovering or rolling back received queue messages makes them available to all consumers on the queue. In previous releases of WebLogic Server, rolled back messages were only available to the session that rolled back the message, until that session was closed.	None required.

---

## Existing Feature 6.0 to 6.1 Functionality Changes

The following table lists the changes in existing feature functionality from WebLogic Server 6.0, and also indicates any code changes that might be required as a result. For additional information pertaining to the JMS Specification's change history, see Chapter 11, "Change History," of [Sun Microsystem's JMS Specification](#)

Category	Description	Code Modification
Connection Factories	<p>For the Acknowledge Policy attribute in the Administration Console, the new default value of <code>All</code> is a work-around to accommodate a change in the <a href="#">JMS Specification</a>. This new default setting represents a change from prior versions of JMS, which internally defaulted to <code>Previous</code>, and which did not appear as an option in the Administration Console.</p> <p>As the message acknowledge policy for the connection factory, the Acknowledge Policy attribute only applies to implementations that use the <code>CLIENT_ACKNOWLEDGE</code> mode for a non-transacted session.</p> <ul style="list-style-type: none"> <li>■ <code>All</code> — acknowledge all messages ever received by a given session, regardless of which message calls the acknowledge method.</li> <li>■ <code>Previous</code> — acknowledge all messages received by a given session, but only up to and including the message that calls the acknowledge method.</li> </ul> <p>For more information on message acknowledge modes, refer to "Non-transacted Session" on page 2-10.</p> <p><b>Note:</b> For connection factories used by MDBs (message-driven beans), always set the Acknowledge Policy field to <code>Previous</code>. Although the default MDB connection factory already does this, foreign connection factories may not.</p>	<p>If you want to acknowledge only previously received messages, up to and including the message that calls the acknowledge method, change the default Acknowledge Policy setting from <code>All</code> to <code>Previous</code> via the <a href="#">JMS Connection Factories</a> tab in the Administration Console.</p>

Category	Description	Code Modification
Destinations	In WLS 6.0, the JMS documentation correctly specifies values of <code>default</code> , <code>true</code> , and <code>false</code> for the <code>StoreEnabled</code> attribute of the <code>JMSDestinationMBean</code> , even though the software allowed for mixed case characters. version 6.1 or later, however, requires all lowercase characters for the <code>StoreEnabled</code> settings.	None required. This is a configuration requirement. For more information, see “ <a href="#">JMS Templates</a> ” in the <i>Administration Console</i> Online Help.

---

# Porting Existing Applications

This release of WebLogic Server supports [Sun Microsystem’s JMS Specification](#). In order to use your existing JMS applications, you must first confirm your version of WebLogic server, and then perform the appropriate porting procedures provided in this section.

- [Porting Steps for 4.5 and 5.1 Applications to 6.x](#)
- [Porting Steps for 6.0 Applications to 6.1](#)
- [Porting Steps for 6.x Applications to 7.0](#)

## Before You Begin

Before beginning the porting procedure, you should check the following list to confirm whether porting is support for your version of WebLogic Server JMS, and to find out whether special porting rules apply to that release:

- Weblogic Server 4.5.1 — Porting is supported *only* for SP15. Customers running all service packs should contact BEA Support.
- Weblogic Server 5.1 — Customers running SP07 or SP08 should contact BEA Support before porting existing JDBC stores to version 7.0.

- In order to port object messages, the object classes need to be in the Weblogic Server 7.0 server `CLASSPATH`.
- For destinations that are not configured in Weblogic Server 7.0, the ported messages will be dropped and the event will be logged.
- WebLogic Server 6.x — All applications are supported in version 7.0. However, if you want your applications to take advantage of the new highly available JMS features, you will need to configure your existing physical destinations (queues and topics) to be part of a single distributed destination set. For more information, see “[Using Distributed Destinations](#)” in *Programming WebLogic JMS*.

## Porting Steps for 4.5 and 5.1 Applications to 6.x

Before you can use an existing WebLogic JMS 6.x application, you must port the WebLogic Server versions 4.5 and 5.1 configuration and message data as follows:

1. Properly shut down the old version of WebLogic Server before beginning the porting process.

**Warning:** Abruptly stopping the old version of WebLogic Server while messaging is still in process may cause problems during porting. Processing should be inactive before shutting down the old server and beginning the porting to WebLogic Server 6.x.

2. Upgrade the WebLogic Server environment, as described in *Installing WebLogic Server*.
3. Ported configuration information using the [configuration conversion facility](#).

During the configuration porting, the following default queue and topic connection factories are enabled:

- `javax.jms.QueueConnectionFactory`
- `javax.jms.TopicConnectionFactory`
- `weblogic.jms.ConnectionFactory`

The first two connection factories are deprecated, but they are still defined and usable for backwards compatibility. For information on the new default connection factory, see the table “[Existing Feature 5.1 to 6.0 Functionality Changes](#)” on page 6-2.

The JMS administrator will need to review the resulting configuration to ensure that the conversion meets the needs of the application. In this case, all of the JMS attributes will be mapped to a single node, as in version 5.1.

**Note:** In versions 6.0 or later, JMS queues are defined during configuration, and no longer saved within database tables. Message data and durable subscriptions are stored either in two JDBC tables or via a directory within the file system.

4. Prepare for automatic porting of existing JDBC database stores.
  - a. Make a backup of the existing JDBC database.
  - b. Ensure that the ported configuration information (see step 2) contains a JDBC database store with exactly the same attributes as the existing store, and that the new JMS servers that use the store define the same destinations and corresponding destination attributes as the existing JMS servers.
  - c. If the new JDBC database store already exists, ensure that it is empty.

The new JDBC database store will be created during the automatic porting, if required.
  - d. Ensure that there is twice the amount of disk space required by the JDBC database store available on the system.

Both the existing and new database information will exist on disk while the porting is performed, doubling the space requirements. Once porting is complete, you can delete the old JDBC database stores, as described in [“Deleting JDBC Database Stores”](#) on page 6-12.
5. Update any existing code, as required, to reflect the feature functionality changes described in [“Existing Feature 5.1 to 6.0 Functionality Changes”](#) on page 6-2.
6. Start up the WebLogic Server and the existing JDBC database stores will be ported automatically.

**Note:** If the automatic porting fails for any reason, the automatic porting will be re-attempted the next time the WebLogic Server boots.

## Porting Steps for 6.0 Applications to 6.1

Before you can use an existing WebLogic JMS 6.x application, you must port the WebLogic Server 6.0 configuration and message data as follows

1. Check the connection factory configuration for version 6.0. You may need to modify programs that call the version 6.1 default connection factory so that they load one of the following connection factories:
  - One of the version 6.0 default connection factories.
  - A custom connection factory.
2. Properly shut down the version 6.0 WebLogic Server before beginning the porting process.

**Warning:** Abruptly stopping the old version of WebLogic Server while messaging is still in process may cause problems during porting. Processing should be inactive before shutting down the old server and beginning the porting to WebLogic Server 6.x.

3. Upgrade the WebLogic Server environment, as described in [Installing WebLogic Server](#).
4. Update any existing code, as required, to reflect the feature functionality changes described in “[Existing Feature 5.1 to 6.0 Functionality Changes](#)” on page 6-2.

**Warning:** Before starting the version 6.1 WebLogic Server, you may want to backup your version 6.0 stores. This is because version 6.0 servers *cannot* use 6.1 stores, and any attempts to do so may cause data corruption.

5. Start the version 6.1 WebLogic Server. This server will continue to use the previous version 6.0 stores.

## Porting Steps for 6.x Applications to 7.0

All WebLogic JMS 6.x applications are supported in version 7.0. However, if you want your applications to take advantage of the new highly available JMS features, you will need to configure your existing physical destinations (queues and topics) to be part of a single distributed destination set.

For more information on using JMS distributed destinations, see “[Using Distributed Destinations](#)” in *Programming WebLogic JMS*.

# Deleting JDBC Database Stores

Once the porting is complete, the old JDBC database tables should be removed using the `utils.Schema` utility, described in detail in [Appendix B, “JDBC Database Utility.”](#)

During porting, a DDL file is generated and stored in the local working directory. The DDL file is named `drop_<jmsServerName>_oldtables.ddl`, where `<jmsServerName>` specifies the name of the JMS server. To delete the JDBC database stores, you pass the resulting DDL file as an argument to the `utils.Schema` utility.

For example, to delete the old JDBC database store from a JMS server named *MyJMSServer*, run the following command:

```
java utils.Schema jdbc:weblogic:oracle weblogic.jdbc.oci.Driver -s
server -u user1 -p foobar -verbose drop_MyJMSServer_oldtables.ddl
```

For more information on the `utils.Schema` utility, see [Appendix B, “JDBC Database Utility.”](#)

# A Configuration Checklists

The following sections provide monitoring checklists for various WebLogic JMS features:

- [Server Clusters](#)
- [JTA User Transactions](#)
- [JMS Transactions](#)
- [Message Delivery](#)
- [Asynchronous Message Delivery](#)
- [Persistent Messages](#)
- [Concurrent Message Processing](#)
- [Multicasting](#)
- [Durable Subscriptions](#)
- [Destination Sort Order](#)
- [Temporary Destinations](#)
- [Thresholds and Quotas](#)

For more information on setting the configuration attributes, refer to the *Administration Guide*. For detailed descriptions of each of the configuration attributes, refer to the *Administration Console Online Help*.

## Server Clusters

To support server clusters, configure the following:

- WebLogic Server targets under the Targets tab on the Connection Factories node
- WebLogic Server targets under the Targets tab on the JMS Servers node

## JTA User Transactions

To support JTA user transactions, configure the following:

- Connection factory JTA user transaction mode by selecting the User Transactions Enabled check box under the Configuration—Transactions tab on the Connection Factories node

## JMS Transactions

To support JMS transacted sessions, configure the following:

- Connection factory transaction timeout value by setting the Transaction Timeout attribute under the Configuration—Transactions tab on the Connection Factories node
- Session pool transaction mode by selecting the Transacted check box under the Configuration tab on the Session Pools node

# Message Delivery

To define message delivery attributes, configure the following:

- Connection factory priority, time-to-live, time-to-deliver, and delivery mode attributes under the Configuration—General tab on the Connection Factories node
- Destination priority, time-to-live, time-to-deliver, and delivery mode override attributes under the Configuration—Overrides tab on the Destinations node
- Destination redelivery delay, redelivery limit, and error destination attributes under the Configuration—Redelivery tab on the Destinations node

**Note:** These settings can also be set dynamically by the message producer when sending a message or using the set methods, as described in [“Sending Messages” on page 4-22](#).

The destination configuration attributes take precedence over all other settings.

## Asynchronous Message Delivery

To define the maximum number of messages that may exist for an asynchronous session and that have not yet been passed to the message listener, configure the following:

- Message maximum attribute under the Configuration—General tab on the Connection Factories node

## Persistent Messages

**Note:** Topic destinations are persistent if, and only if they have durable subscriptions. For more information about durable subscriptions, see [“Setting Up Durable Subscriptions”](#) on page 4-53.

To support persistent messaging, configure the following:

- Create a file or JDBC store using the Stores node
- JMS server backing store by setting the Store attribute under the Configuration—General tab on the JMS Servers node
  - Note:** No two JMS servers can use the same backing store.
- Default message delivery mode by setting one of the following attributes to `PERSISTENT` or `NON_PERSISTENT`:
  - Default Delivery Mode attribute under the Configurations—General tab on the Connection Factories node
  - Delivery Mode Override attribute under the Configurations—Overrides tab on the Destination node

**Note:** You can also specify persistent as the delivery mode when sending messages, as described in [“Sending Messages”](#) on page 4-22.

## Concurrent Message Processing

To support concurrent message processing, configure the following:

- Server session pool attributes under the Configuration tab on the Session Pools node
- Connection consumer attributes under the Configuration tab on the Connection Consumers node

**Note:** Server session pool factories, used for concurrent message processing, are not configurable. WebLogic JMS defines one `ServerSessionPoolFactory` object, by default: `weblogic.jms.ServerSessionPoolFactory:<name>`, where `<name>` specifies the name of the JMS server on which the session pool is created. For more information about using server session pool factories, refer to [“Defining Server Session Pools” on page 4-71](#).

## Multicasting

**Note:** Multicasting applies to topics only.

To configure multicasting on a topic, configure the following:

- Multicast address, multicast port, and multicast time-to-live (TTL) under the Configuration—Multicast tab on the Destination node
- Maximum number of outstanding messages by setting the Messages Maximum attribute under the Configuration—General tab on the Connection Factories node
- Overrun policy used when the number of outstanding messages reaches the Messages Maximum value by setting the Overrun Policy attribute under the Configuration—General tab on the Connection Factories node

## Durable Subscriptions

To support durable topic subscriptions, configure the following:

- Create a JMS file store or JMS JDBC persistent store using the Stores node.
- Target the configured store to your JMS server by selecting it from the Store attribute’s drop-down list on the JMS Server → Configuration → General tab.

**Note:** No two JMS servers can use the same persistent store.

- Set the client ID on either a JMS connection factory or in the topic connection, as follows:

- One method is to configure a connection factory with the client identifier (client ID). For WebLogic JMS, this means configuring a separate connection factory instance for each client ID. Then, on each connection factory, specify a unique client ID for clients with durable subscriptions using the Client ID attribute on the Connection Factory → Configuration → General tab. Applications then look up their own topic connection factories in JNDI and use them to create connections containing their own client IDs.
- The other “preferred” method is for an application that can set its client ID in the topic connection after the connection is created by calling a connection method, as described in [“Setting Up Durable Subscriptions” on page 4-53](#). By using this alternative approach, you can use the default connection factory (if it is acceptable for your application) and avoid the need to modify the configuration information.

## Destination Sort Order

To support destination sort order, configure the following:

- Key attributes under the Configuration tab on Destination Keys node
- Destination Keys under Configuration—General tab on Destinations node

## Temporary Destinations

To support temporary destinations (queue or topic), configure the following:

- A JMS template for the JMS server (in the same domain) under the Configuration—General tab on the Templates node
- A JMS template to be used by the JMS server for temporary destinations by setting the Temporary Template attribute for the JMS server under the Configuration—General tab on the JMS Servers node

# Thresholds and Quotas

To configure thresholds and quotas, configure the following:

- Message and byte thresholds and quotas (maximum number, and high and low thresholds) under the Configurations—Thresholds tab on the JMS Server node
- Message and byte thresholds and quotas (maximum number, and high and low thresholds) under the Configurations—Thresholds tab on the Destination node
- Maximum number of sessions that can be retrieved from a session pool by setting the Sessions Maximum attribute under the Configurations tab on the Session Pools node
- Maximum number of messages that can be accumulated by a connection consumer by setting the Messages Maximum attribute under the Configuration tab of the Consumers node



# B JDBC Database Utility

The following sections describe JDBC database stores for WebLogic JMS, and how to use the JDBC database utility to regenerate existing JDBC database stores:

- [Overview](#)
- [About JMS Tables](#)
- [Regenerating JDBC Database Stores](#)

## Overview

The JDBC `utils.Schema` utility allows you to regenerate new JDBC database stores by deleting the existing versions. Running this utility is usually not necessary, since JMS automatically creates these stores for you. However, if your existing JDBC database stores somehow become corrupted, you can regenerate them using the `utils.Schema` utility.

**Caution:** Use caution when running the `utils.Schema` command as it will delete all existing database tables and then recreate new ones.

## About JMS Tables

The JMS database contains two system tables that are generated automatically and are used internally by JMS, as follows:

- `<prefix>JMSStore`
- `<prefix>JMSState`

The prefix name uniquely identifies JMS tables in the backing store. Specifying unique prefixes allows multiple stores to exist in the same database. The prefix is configured via the Administration Console when configuring the JDBC store. A prefix is prepended to table names when:

- The DBMS requires fully qualified names.
- You must differentiate between JMS tables for two WebLogic servers, enabling multiple tables to be stored on a single DBMS.

The prefix should be specified using the following format, which will result in a valid table name when prepended to the JMS table name:

```
[[catalog.]schema.]prefix
```

**Note:** No two JMS stores should be allowed to use the same database tables, as this will result in data corruption.

For more information on configuring JDBC database stores for WebLogic JMS, see [Managing JMS](#) in the *Administration Guide*.

# Regenerating JDBC Database Stores

The `utils.Schema` utility is a Java program that takes command line arguments to specify the following:

- JDBC driver
- Database connection information
- Name of a file containing the SQL Data Definition Language (DDL) commands (terminated by semicolons) that create the database tables

By convention, the DDL file has a `.ddl` extension. DDL files are provided for Pointbase, Cloudscape, Informix, Sybase, Oracle, MS SQL Server, IBM DB2, and Times Ten databases.

To execute `utils.Schema`, your `CLASSPATH` must contain the `weblogic.jar` file.

Enter the `utils.Schema` command, as follows:

```
java utils.Schema url JDBC_driver [options] DDL_file
```

The following table lists the `utils.Schema` command-line arguments.

Argument	Description
<i>url</i>	Database connection URL. This value must be a colon-separated URL as defined by the JDBC specification.
<i>JDBC_driver</i>	Full package name of the JDBC Driver class.
<i>options</i>	<p>Optional command options.</p> <p>If required by the database, you can specify:</p> <ul style="list-style-type: none"><li>■ The username and password as follows: <code>-u &lt;username&gt; -p &lt;password&gt;</code></li><li>■ The Domain Name Server (DNS) name of the JDBC database server as follows: <code>-s &lt;dbserver&gt;</code></li></ul> <p>You can also specify the <code>-verbose</code> option, which causes <code>utils.Schema</code> to echo SQL commands as they are executed.</p>
<i>DDL_file</i>	<p>The full pathname of a text file containing the SQL commands that you wish to execute. An SQL command can span several lines and is terminated with a semicolon (;). Lines beginning with pound signs (#) are comments.</p> <p>The <code>weblogic/jms/ddl</code> directory within the <code>weblogic.jar</code> file contains JMS DDL files for Pointbase, Cloudscape, Informix, Sybase, Oracle, MS SQL Server, IBM DB2, and Times Ten databases. To use a different database, copy and edit any one of these files.</p> <p>Use the jar utility supplied with the JDK to extract the DDL files to the <code>weblogic/jms/ddl</code> directory, using the following command:</p> <pre>jar xf weblogic.jar weblogic/jms/ddl</pre> <p><b>Note:</b> If you omit the <code>weblogic/jms/ddl</code> parameter, the entire jar file is extracted.</p>

For example, the following command recreates the JMS tables in an Oracle server named `DEMO`, with the username `user1` and password `foobar`:

```
java utils.Schema jdbc:weblogic:oracle:DEMO \  
  weblogic.jdbc.oci.Driver -u user1 -p foobar -verbose \  
  weblogic/jms/ddl/jms_oracle.ddl
```

With the Pointbase demo database that is shipped with WebLogic Server, no username or password is required. However, you must follow this procedure to create the JMS tables in a Pointbase server:

1. Set the WLS samples environment:

```
%SAMPLES_HOME%\server\config\examples\setExamplesEnv.cmd
```

2. Change to the `%WL_HOME%\server\lib\directory`, and then extract the `jms_pointbase.ddl` file from the `weblogic.jar` file to the current directory.
3. Execute the following command to create the JMS tables:

```
java utils.Schema jdbc:pointbase:server://localhost/demo  
  com.pointbase.jdbc.jdbcUniversalDriver  
  -u examples -p examples -verbose jms_pointbase.ddl
```

The Pointbase JDBC URL specifies the demo database, which is included with the WebLogic JMS samples. For the samples, the JMS tables have already been created in this database.

4. Start the Pointbase server and open the Pointbase console.

For detailed information on using the Pointbase Server console to monitor and manipulate the JMS tables, see the `Pointbase.html` file in the `WL_HOME\samples\server\src\examples` directory, where `WL_HOME` is the top-level directory of your WebLogic Platform installation.

---

# Index

## A

- Acknowledge message 4-33
- Acknowledge modes 2-10
- Anonymous producer 4-25, 4-26
- Application development flow
  - acknowledging received messages 4-33
  - importing required packages 4-3
  - receiving messages 4-29
  - releasing object resources 4-34
  - sending messages 4-22
  - setting up 4-4
  - steps 4-2
- Application setup
  - creating a connection 4-7
  - creating a session 4-8
  - creating message consumers 4-11
  - creating message object 4-13
  - creating message producers 4-11
  - example
    - PTP 4-16
    - Pub/sub 4-19
  - looking up connection factory 4-6
  - looking up destination 4-9
  - receiving messages asynchronously 4-15
  - registering asynchronous message listener 4-15
  - starting the connection 4-16
  - steps 4-4
- Asynchronous message, receiving 4-15, 4-30

## B

- Bytes message
  - creating 4-14

## C

- Client ID
  - defining 4-55
  - displaying 4-56
- Close
  - connection 4-47
  - session 4-49
- Clusters
  - configuration checklist A-2
  - configuring 3-3
- Concurrent processing 4-71
- Configuration
  - checklists A-1
  - clustered JMS 3-3
  - JMS 3-1
  - migratable targets 3-6
- Connection
  - closing 4-47
  - creating 4-7
  - definition of 2-8
  - exception listener 4-44
  - managing 4-44
  - metadata 4-45
  - starting 4-16, 4-46
  - stopping 4-46
- Connection consumer

---

- definition of 2-23
- queue 4-76
- topic 4-76

Connection factory

- definition of 2-6
- looking up 4-6

customer support contact information xiii

## D

Delivery mode 4-24, 4-25, 4-27

Delivery time

- overriding
  - on destinations 4-40
  - relative time-to-deliver 4-40
  - schedule interface 4-43
  - scheduled time-to-deliver syntax 4-41
- scheduling overview 4-38
- setting on messages 4-39
- setting on producer 4-39

Destination

- creating dynamically 4-49
- definition of 2-13
- looking up 4-9
- sort order 4-29
- temporary 4-52

Destination, distributed

- definition of 2-14

documentation, where to find it xii

Durable subscription

- client ID 4-55
- creating 4-56
- deleting 4-57
- modifying 4-57
- setting up 4-53

## E

Error destination for undelivered messages 4-37

Error recovery

- connection 4-44
- session 4-48

Examples

- browse queue 4-66
- closing resources 4-35
- JMS and EJB in JTA user transaction 5-9
- message filtering 4-68
- multicast session 4-86
- receiving messages synchronously
  - PTP 4-32
  - Pub/sub 4-32
- sending messages
  - PTP 4-28
  - Pub/sub 4-28
- server session pool
  - PTP 4-77
  - Pub/sub 4-79
- setting message header field 4-61
- setting up
  - PTP 4-16
  - Pub/sub 4-19

Exception listener

- connection 4-44
- session 4-48

Existing feature functionality changes 6-1

## F

Failover procedures 3-11

Failure, server 3-11

Filter message

- definition 4-66
- example 4-68
- SQL statement 4-67
- XML selector 4-68

## H

Header fields

- browsing 4-65

---

- definition of 2-16
- displaying 4-59
- setting 4-59

## **J**

### JDBC store

- automatic porting 6-10
- database utility B-1

### JMS

- architecture 1-3
  - clustering features 1-5
  - major components 1-4
- classes 2-5
- configuring 3-1
- configuring clusters 3-3
- configuring migratable targets 3-6
- definition 1-1
- existing feature functionality changes
  - 6-1
- features 1-3
- monitoring 3-10
- tuning 3-9

### JMS transacted sessions

- committing or rolling back 5-4
- configuration checklist A-2
- creating 5-4
- displaying 5-4
- executing operations 5-4

### JMSCorrelationID header field

- definition of 2-17
- displaying 4-59
- setting 4-59

### JMSDeliveryMode header field

- definition of 2-18
- displaying 4-59, 4-60

### JMSDeliveryTime header field

- definition of 2-18
- displaying 4-60

### JMSDestination header field

- definition of 2-18

- displaying 4-59

### JMSExpiration header field

- definition of 2-18

### JMSHelper class methods 4-50

### JMSMessageID header field

- definition of 2-19

- displaying 4-60

### JMSPriority header field

- definition of 2-19

- displaying 4-60

### JMSRedelivered header field

- definition of 2-19

- displaying 4-60

### JMSReplyTo header field

- definition of 2-20

- displaying 4-61

- setting 4-61

### JMSTimestamp header field

- definition of 2-20

- displaying 4-61

- setting 4-61

### JMSType header field

- definition of 2-20

- displaying 4-61

- setting 4-61

### JTA user transaction

- committing or rolling back 5-8

- configuration checklist A-2

- creating non-transacted session 5-7

- example 5-9

- looking up user transaction in JNDI 5-7

- performing desired operations 5-8

- starting 5-7

## **M**

### Map message

- creating 4-14

### Message

- acknowledging 4-33

- body 2-21

---

- creating object 4-13, 4-22
- defining content 4-22
- definition 1-1
- definition of 2-16
- delivery
  - configuration checklists A-3
  - mode 4-24, 4-25, 4-27
  - times, setting 4-38
- filtering
  - definition 4-66
  - SQL message selector 4-67
  - XML message selector 4-68
- header fields
  - browsing 4-65
  - definition of 2-16
  - displaying 4-59
  - setting 4-59
- managing
  - rolled back and recovered 4-35
- persistence
  - configuration checklist A-4
  - definition of 2-4
- priority 4-24, 4-25, 4-27
- property fields
  - browsing 4-65
  - clearing 4-61
  - conversion chart 4-63
  - definition of 2-20
  - displaying 4-61
  - displaying all 4-63
  - setting 4-61
- receiving
  - asynchronous 4-15, 4-30
  - order control 4-29
  - synchronous 4-31
- recovering 4-33
- redelivery delay 4-35
- redelivery limit 4-37
- sending 4-22
- server session pools 4-71
- setting delivery times 4-38
- time-to-deliver 4-27, 4-40, 4-41
- time-to-live 4-24, 4-25, 4-27
- types
  - definition of 2-21
  - displaying 4-62
  - setting 4-14, 4-62
- Message consumer
  - creating 4-11
  - definition of 2-14
- Message driven beans 5-9
- Message listener, registering 4-15
- Message producer
  - creating 4-11
  - creating dynamically 4-26
  - definition of 2-14
- Message selector
  - defining
    - SQL 4-67
    - XML 4-68
  - displaying 4-69
  - example 4-68
- Messaging models
  - point-to-point 2-2
  - publish/subscribe 2-3
- Metadata, connection 4-45
- Migratable targets
  - configuring 3-6
- Monitor JMS 3-10
- Multicast session
  - creating 4-83
  - creating topic subscriber 4-83
  - definition 4-81
  - dynamically configuring 4-85
  - example 4-86
  - messages maximum 4-85
  - overrun policy 4-85
  - prerequisites 4-82
  - setting up message listener 4-84

---

## **N**

Non-durable subscription 4-54

## **O**

Object message  
  creating 4-14

Overriding

  delivery time  
    overview 4-40  
    relative time-to-deliver 4-40  
    schedule interface 4-43  
    scheduled time-to-deliver syntax  
      4-41  
  redelivery delay 4-36

## **P**

Packages, required 4-3

Persistent message  
  configuration checklist A-4  
  definition of 2-4

Point-to-point messaging  
  definition of 2-2  
  example  
    receiving messages synchronously  
      4-32  
    sending messages 4-28  
    server session pool 4-79  
    setting up application 4-16

Porting procedures 6-8  
  steps for 4.5 and 5.1 applications to 6.x  
    6-9  
  steps for 6.0 applications to 6.1 6-11  
  steps for 6.x applications to 7.0 6-11

printing product documentation xii

Priority, message 4-24, 4-25, 4-27

Property fields

  browsing 4-65  
  clearing 4-61  
  conversion chart 4-63

  displaying 4-61

  displaying all 4-63

  setting 4-61

Publish/subscribe messaging

  definition of 2-3

  example

    receiving messages synchronously  
      4-32

    sending messages 4-28

    setting up application 4-19

## **Q**

Queue

  creating 4-10  
  creating dynamically 4-49  
  definition of 2-13  
  displaying 4-10, 4-12  
  temporary  
    creating 4-52  
    definition of 2-13  
    deleting 4-53

Queue connection

  creating 4-7  
  definition of 2-9

Queue connection factory

  creating queue connection 4-7  
  definition of 2-8  
  looking up 4-6

Queue receiver

  creating 4-11  
  definition of 2-15  
  receiving messages 4-31

Queue sender

  creating 4-11  
  definition of 2-15  
  sending message 4-23

Queue session

  creating 4-8  
  definition of 2-10

---

## R

- Receive message
    - asynchronous 4-15, 4-30
    - order 4-29
    - synchronous 4-31
  - Recover from system failure 3-11
  - Recover message 4-33, 4-35
  - Redeliver message 4-33
  - Redelivery delay
    - overriding on destination 4-36
    - overview 4-35
    - setting for messages 4-36
  - Redelivery limit 4-27
    - configuring error destination 4-37
    - configuring limit 4-37
    - overview 4-37
  - Release object resources 4-34
  - Request/response, support of 2-17
  - Resources, releasing 4-34
  - Rolled back messages
    - managing 4-35
      - redelivery delay 4-35
      - redelivery limit 4-37
- ## S
- Send messages 4-22
  - Server failure recovery 3-11
  - Server session
    - definition of 2-23
    - retrieving 4-75
  - Server session pool
    - creating
      - queue connection consumers 4-74
      - topic connection consumers 4-75
    - definition of 2-22
    - setting up 4-71
  - Server session pool factory
    - creating a server session pool 4-74
    - definition of 2-22
    - looking up 4-73

## Session

- acknowledge modes 2-10
  - closing 4-49
  - creating 4-8
  - definition of 2-9
  - exception listener 4-48
  - managing 4-47
  - non-transacted 2-10
  - transacted 2-12
- ## SQL message selectors 4-67
- ## Start connection 4-16, 4-46
- ## Stop connection 4-46
- ## Stream message
- creating 4-14
- ## support
- technical xiii
- ## Synchronous receive 4-31

## T

- ## Temporary destination
- configuring server A-6
  - creating
    - queue 4-52
    - topic 4-53
  - deleting 4-53
- ## Temporary queue
- creating 4-52
  - definition of 2-13
  - deleting 4-53
- ## Temporary topic
- creating 4-53
  - definition of 2-13
  - deleting 4-53
- ## Time-to-deliver 4-27, 4-40
- ## Time-to-live 4-24, 4-25, 4-27, 4-40
- ## Topic
- creating 4-10
  - creating dynamically 4-49
  - definition of 2-13
  - displaying 4-10, 4-13

---

- displaying NoLocal variable 4-13
- JMSHelper class methods 4-50
- temporary
  - creating 4-53
  - definition of 2-13
  - deleting 4-53
- Topic connection
  - creating 4-7
  - definition of 2-9
- Topic connection factory
  - creating topic connection 4-7
  - definition of 2-8
  - looking up 4-6
- Topic publisher
  - creating 4-12
  - definition of 2-15
  - sending messages 4-25
- Topic session
  - creating 4-9
  - definition of 2-10
- Topic subscriber
  - creating 4-12
  - definition of 2-15
  - durable 4-53
- Transactions 5-1
  - JMS transacted sessions. See JMS transacted sessions
  - JTA user transaction. See JTA user transaction
- Tuning JMS 3-9

## **U**

- utils.Schema utility 6-12, B-1

## **X**

- XML message
  - class 2-21
  - creating 4-14
  - selector 4-68